NavisCore IP Navigator Configuration Guide

Ascend Communications, Inc.

Product Code: 80056
Revision 00
September 1998

ASCEND COMMUNICATIONS, INC. END-USER LICENSE AGREEMENT

ASCEND COMMUNICATIONS, INC. IS WILLING TO LICENSE THE ENCLOSED SOFTWARE AND ACCOMPANYING USER DOCUMENTATION (COLLECTIVELY, THE "PROGRAM") TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. PLEASE READ THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT CAREFULLY BEFORE OPENING THE PACKAGE(S) OR USING THE ASCEND SWITCH(ES) CONTAINING THE SOFTWARE, AND BEFORE USING THE ACCOMPANYING USER DOCUMENTATION. OPENING THE PACKAGE(S) OR USING THE ASCEND SWITCH(ES) CONTAINING THE PROGRAM WILL INDICATE YOUR ACCEPTANCE OF THE TERMS OF THIS LICENSE AGREEMENT. IF YOU ARE NOT WILLING TO BE BOUND BY THE TERMS OF THIS LICENSE AGREEMENT, ASCEND IS UNWILLING TO LICENSE THE PROGRAM TO YOU, IN WHICH EVENT YOU SHOULD RETURN THE PROGRAM WITHIN TEN (10) DAYS FROM SHIPMENT TO THE PLACE FROM WHICH IT WAS ACQUIRED, AND YOUR LICENSE FEE WILL BE REFUNDED. THIS LICENSE AGREEMENT REPRESENTS THE ENTIRE AGREEMENT CONCERNING THE PROGRAM BETWEEN YOU AND ASCEND, AND IT SUPERSEDES ANY PRIOR PROPOSAL, REPRESENTATION OR UNDERSTANDING BETWEEN THE PARTIES.

1. License Grant. Ascend hereby grants to you, and you accept, a non-exclusive, non-transferable license to use the computer software, including all patches, error corrections, updates and revisions thereto in machine-readable, object code form only (the "Software"), and the accompanying User Documentation, only as authorized in this License Agreement. The Software may be used only on a single computer owned, leased, or otherwise controlled by you; or in the event of inoperability of that computer, on a backup computer selected by you. You agree that you will not pledge, lease, rent, or share your rights under this License Agreement, and that you will not, without Ascend's prior written consent, assign or transfer your rights hereunder. You agree that you may not modify, reverse assemble, reverse compile, or otherwise translate the Software or permit a third party to do so. You may make one copy of the Software and User Documentation for backup purposes. Any such copies of the Software or the User Documentation shall include Ascend's copyright and other proprietary notices. Except as authorized under this paragraph, no copies of the Program or any portions thereof may be made by you or any person under your authority or control.

2. Ascend's Rights. You agree that the Software and the User Documentation are proprietary, confidential products of Ascend or Ascend's licensor protected under US copyright law and you will use your best efforts to maintain their confidentiality. You further acknowledge and agree that all right, title and interest in and to the Program, including associated intellectual property rights, are and shall remain with Ascend or Ascend's licensor. This License Agreement does not convey to you an interest in or to the Program, but only a limited right of use revocable in accordance with the terms of this License Agreement.

3. License Fees. The license fees paid by you are paid in consideration of the license granted under this License Agreement.

4. Term. This License Agreement is effective upon your opening of the package(s) or use of the switch(es) containing Software and shall continue until terminated. You may terminate this License Agreement at any time by returning the Program and all copies or portions thereof to Ascend. Ascend may terminate this License Agreement upon the breach by you of any term hereof. Upon such termination by Ascend, you agree to return to Ascend the Program and all copies or portions thereof. Termination of this License Agreement shall not prejudice Ascend's rights to damages or any other available remedy.

5. Limited Warranty. Ascend warrants, for your benefit alone, for a period of 90 days from the date of shipment of the Program by Ascend (the "Warranty Period") that the program diskettes in which the Software is contained are free from defects in material and workmanship. Ascend further warrants, for your benefit alone, that during the Warranty Period the Program shall operate substantially in accordance with the User Documentation. If during the Warranty Period, a defect in the Program appears, you may return the Program to the party from which the Program was acquired for either replacement or, if so elected by such party, refund of amounts paid by you under this License Agreement. You agree that the foregoing constitutes your sole and exclusive remedy for breach by Ascend of any warranties made under this Agreement. EXCEPT FOR THE WARRANTIES SET FORTH ABOVE, THE PROGRAM IS LICENSED "AS IS", AND ASCEND DISCLAIMS ANY AND ALL OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTIES OF NONINFRINGEMENT.

6. Limitation of Liability. Ascend's cumulative liability to you or any other party for any loss or damages resulting from any claims, demands, or actions arising out of or relating to this License Agreement shall not exceed the greater of: (i) ten thousand US dollars ($10,000) or (ii) the total license fee paid to Ascend for the use of the Program. In no event shall Ascend be liable for any indirect, incidental, consequential, special, punitive or exemplary damages or lost profits, even if Ascend has been advised of the possibility of such damages.

7. Proprietary Rights Indemnification. Ascend shall at its expense defend you against and, subject to the limitations set forth elsewhere herein, pay all costs and damages made in settlement or awarded against you resulting from a claim that the Program as supplied by Ascend infringes a United States copyright or a United States patent, or misappropriates a United States trade secret, provided that you: (a) provide prompt written notice of any such claim, (b) allow Ascend to direct the defense and settlement of the claim, and (c) provide Ascend with the authority, information, and assistance that Ascend deems reasonably necessary for the defense and settlement of the claim. You shall not consent to any judgment or decree or do any other act in compromise of any such claim without first obtaining Ascend's written consent. In any action based on such a claim, Ascend may, at its sole option, either: (1) obtain for you the right to continue using the Program, (2) replace or modify the Program to avoid the claim, or (3) if neither (1) nor (2) can reasonably be effected by Ascend, terminate the license granted hereunder and give you a prorata refund of the license fee paid for such Program, calculated on the basis of straight-line depreciation over a five-year useful life. Notwithstanding the preceding sentence, Ascend will have no liability for any infringement or misappropriation claim of any kind if such claim is based on: (i) the use of other than the current unaltered release of the Program and Ascend has provided or offers to provide such release to you for its then current license fee, or (ii) use or combination of the Program with programs or data not supplied or approved by Ascend to the extent such use or combination caused the claim.

8. Export Control. You agree not to export or disclose to anyone except a United States national any portion of the Program supplied by Ascend without first obtaining the required permits or licenses to do so from the US Office of Export Administration, and any other appropriate government agency.

9. Governing Law. This License Agreement shall be construed and governed in accordance with the laws and under the jurisdiction of the Commonwealth of Massachusetts, USA. Any dispute arising out of this Agreement shall be referred to an arbitration proceeding in Boston, Massachusetts, USA by the American Arbitration Association.

10. Miscellaneous. If any action is brought by either party to this License Agreement against the other party regarding the subject matter hereof, the prevailing party shall be entitled to recover, in addition to any other relief granted, reasonable attorneys' fees and expenses of arbitration. Should any term of this License Agreement be declared void or unenforceable by any court of competent jurisdiction, such declaration shall have no effect on the remaining terms hereof. The failure of either party to enforce any rights granted hereunder or to take action against the other party in the event of any breach hereunder shall not be deemed a waiver by that party as to subsequent enforcement of rights or subsequent actions in the event of future breaches.

Contents

Chapter 2 Configuring Ethernet Logical Ports

Chapter 3 Configuring IP Logical Ports and IP Servers

Chapter 4 Configuring IP Packet Filters

Chapter 5 Provisioning IP Quality of Service

Chapter 6 Configuring Static ARP Entries

Chapter 7 Configuring RIP

Chapter 8 Configuring BGP Parameters

Chapter 9 Configuring OSPF Parameters

List of Figures

List of Tables

About This Guide

The NavisCore IP Navigator Configuration Guide is a task-oriented guide that describes, step-by-step, the process for configuring an IP interface. This guide is intended for users who will be accessing the NavisCore NMS to configure IP interfaces in a network.

What You Need to Know

As a reader of this guide, you should be familiar with the UNIX operating system and HP OpenView. The system administrator should be familiar with relational database software to properly maintain Sybase, the database used by NavisCore.

This guide assumes you have already installed the Ascend switch hardware, using one of the following guides:

- *B-STDX 8000/9000 Hardware Installation Guide*
- *CBX 500 Hardware Installation Guide*

and you have installed the NMS software, using the *Network Management Station Installation Guide*.

You should also possess a working knowledge of Internet routing architectures and IP routing procedures. In addition, you should be familiar with the configuration process for a Frame Relay or ATM logical port. Configuration instructions for an Ethernet logical port are included in this guide. See page 1-8 for a list of the logical ports that support IP routing on the B-STDX 8000/9000 and on the CBX 500.

Reading Path

This section describes all of the documents that support the NavisCore NMS and switch software. The documents are grouped as follows:

- NMS Documentation
- Switch Software Documentation
- Third Party Documentation

NMS Documentation

Read the following documents to install and operate NavisCore Release 4.0. Be sure to review the NavisCore Customer Software Release Notice for any changes not included in these guides.

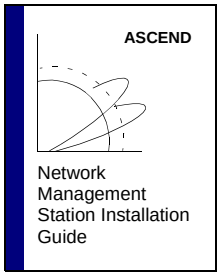

This guide describes prerequisite tasks, hardware and software requirements, and instructions for installing Solaris, HP OpenView, and NavisCore on the NMS.

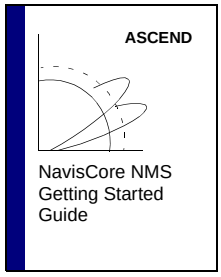

This guide describes how to configure and manage NavisCore, network maps, and Ascend switches.

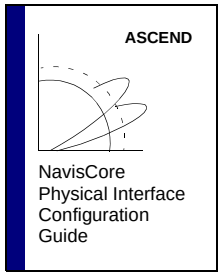

This guide describes how to configure processor and I/O modules on Ascend switches.

Switch Software Documentation

Read the following documents to configure switch software for B-STDX Release 6.0, CBX Release 3.0, and GX Release 1.0. Be sure to review the appropriate switch software customer software release notice(s) for any changes.

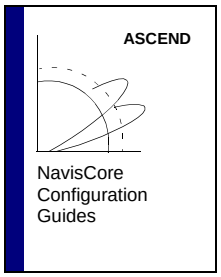

These guides describe how to configure WAN services on the STDX, B-STDX, CBX, and GX switch platforms:

- *NavisCore Frame Relay Configuration Guide*
- *NavisCore ATM Configuration Guide*

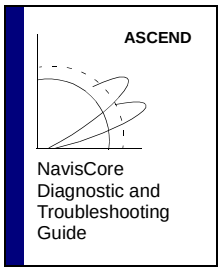

This guide describes how to diagnose and troubleshoot your NavisCore switch network.

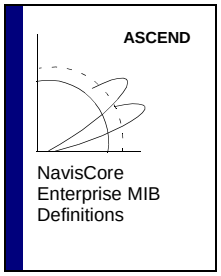

This document gives a brief overview of SNMP and describes the NavisCore Enterprise MIB definitions.

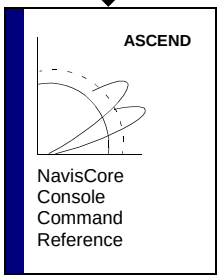

This reference lists and describes the NavisCore switch console commands.

Third Party Documentation

- Halabi, Bassam. *Internet Routing Architectures.* New Riders Publishing, 1997.
- Moy, John T. *OSPF Anatomy of an Internet Routing Protocol.* Addison-Wesley Publishing Group, 1998.
- Reynolds, J. *Postel.RFC 1700 ASSIGNED NUMBERS.* October 1994.

How to Use this Guide

The following table highlights the chapters in this guide.

Read	To Learn About
Chapter 1	An overview of the product.
Chapter 2	How to configure logical ports on the following cards: • 4-Port Ethernet Card on the CBX • 2-Port Ethernet Card on the B-STDX
Chapter 3	How to configure: • IP logical ports on B-STDX and CBX switches. IP logical ports are ports that support IP routing. • IP Server logical ports on the CBX. IP Server logical ports provide a method of accepting or transmitting IP traffic on a cell-based card.
Chapter 4	How to configure and assign IP packet filters.
Chapter 5	How to provision IP Quality of Service (QoS).
Chapter 6	How to define static ARP entries.
Chapter 7	How to configure Routing Information Protocol (RIP) parameters on an IP logical port.
Chapter 8	How to configure Border Gateway Protocol (BGP) parameters including: • BGP switch parameters • BGP neighbors • BGP aggregates
Chapter 9	How to configure Open Shortest Path First (OSPF) parameters for IP services.

Read	To Learn About
Chapter 10	How to manually configure static routes.
Chapter 11	How to create a route map. The purpose of a route map is to control and modify routing information and to define the parameters that your system uses to redistribute routes between routing domains.
Chapter 12	How IP Navigator uses Multipoint-to-Point Tunnels (MPTs) as a means of forwarding IP traffic over switched paths through the Ascend network.
Appendix A	The Upload PRAM function.

What's New in This Release?

This guide describes the following new product features:

New Features/Functions	Enables You To	Described in Chapter
Ethernet Cards	Configure Ethernet logical ports on the following cards: • 2-port Ethernet card on the B-STDX • 4-port Ethernet card on the CBX 500	Chapter 2
IP logical ports on CBX 500 switches	Configure IP logical ports on CBX 500 switches. The following CBX 500 cards support IP logical ports: • 6-port DS3 Frame card • 4-port Ethernet card	Chapter 3
IP Server functions	Configure logical ports as IP Server logical ports on either of the following CBX 500 cards: • 6-port DS3 Frame card • 4-port Ethernet card The purpose of an IP Server logical port is to provide a method of accepting or transmitting IP traffic on a cell-based card.	Chapter 3
Area border routers	Have switches/routers with links to more than one area, or links between an area and a backbone.	Chapter 9
Multiple OSPF Area Support	Use a hierarchical architecture in large networks. This architecture can improve the performance of route look-ups and reduce the routing table size.	Chapter 9
MPT point-to-point connections for routed paths.	Create user-defined circuits for IP traffic between two switches so that you can provide more bandwidth between two nodes.	Chapter 12

Conventions

This guide uses the following conventions to emphasize certain information, such as user input, screen prompts and output, and menu selections. For example:

Convention	Indicates	Example
`Courier Bold`	User input on a separate line.	**`eject cdrom`**
`Courier`	Screen or system output.	`Please wait...`
[*bold italics*]	Variable parameters to enter.	**[*your IP address*]**
`<Return>`	Press Return or Enter.	`<Return>`
Boldface	User input in text.	Type **cd install** and. . .
Menu ⟹ Option	Select an option from the menu.	CascadeView ⟹ Logon
Boxes surrounding text	Notes and warnings.	See examples below.
Italics	Book titles, new terms, filenames, directories, and emphasized text.	*Network Management Station Installation Guide*

Notes provide additional information or helpful suggestions that may apply to the subject text.

Cautions notify the reader to proceed carefully to avoid possible equipment damage or data loss.

Warnings notify the reader to proceed carefully to avoid possible personal injury.

Related Documents

This section lists the related Ascend documentation that may be helpful to read.

- *Network Management Station Installation Guide* (Product Code: 80014)
- *NavisCore NMS Getting Started Guide* (Product Code: 80070)
- *NavisCore Physical Interface Configuration Guide* (Product Code: 80080)
- *NavisCore Frame Relay Configuration Guide* (Product Code: 80071)
- *NavisCore ATM Configuration Guide* (Product Code: 80072)
- *NavisCore Diagnostic and Troubleshooting Guide* (Product Code: 80074)
- *NavisCore Console Command Reference* (Product Code: 80075)
- *B-STDX 8000/9000 Hardware Installation Guide* (Product Code: 80005)
- *CBX 500 Hardware Installation Guide* (Product Code: 80011)

Customer Comments

Customer comments are welcome. Please respond in one of the following ways:

- Fill out the Customer Comment Form located at the back of this guide and return it to us.
- E-mail your comments to cspubs@ascend.com.
- FAX your comments to 978-692-1510, attention Technical Publications.
- Open a case in CaseView for documentation.

Customer Support

To obtain release notes, technical tips, or support, access the Ascend FTP Server or contact the Technical Assistance Center (TAC) at:

- 1-800-DIAL-WAN or 1-978-952-7299 (U.S. and Canada)
- 0-800-96-2229 (U.K.)
- 1-978-952-7299 (all other areas)

Common Acronyms

The following table lists and describes some of the acronyms used throughout this guide.

Acronym	Description
ARP	Address Resolution Protocol
AS	Autonomous System
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
CIDR	Classless Inter-Domain Routing
DVMRP	Distance Vector Multicast Routing Protocol
FTP	File Transfer Protocol
ICMP	Internet Control Message Protocol
IFMP	Ipsilon Flow Management Protocol (RFC 1953)
IGMP	Internet Group Multicast Protocol
InARP	Inverse Address Resolution Protocol
IP	Internet Protocol
MOSPF	Multicast Open Shortest Path First
MPOA	Multi Protocol Over ATM
MPT	Multipoint-to-Point Tunneling
NBMA	Non-Broadcast Multi-Access
NHRP	Next Hop Resolution Protocol
OSPF	Open Shortest Path First
PIM	Protocol Independant Multicast
PNNI	Private Network to Network Interface
PPP	Point-to-Point
PVC	Permanent Virtual Circuit
RARP	Reverse Address Resolution Protocol
RIP	Routing Internet Protocol

Acronym	Description
SVC	Switched Virtual Circuit
TCP	Transmission Control Protocol
VCI	Virtual Circuit Identifier
VPI	Virtual Path Identifier
VPN	Virtual Private Network
VNN	Virtual Network Navigator

1

Overview

This chapter provides an overview of Ascend's IP switching technology, called IP Navigator™, and describes how Ascend uses IP Navigator to implement the TCP/IP protocol suite into its multiservice switching platforms.

About IP Switching

IP Switching technology allows Ascend's multiservice WAN switching platforms to assume the characteristics and role of an IP router. The main difference between Ascend's IP switching and traditional IP routing is that in the core of the Ascend network, IP packets are switched instead of routed. In other words, instead of examining IP headers at each hop, Ascend switches examine the IP header only at the ingress and egress ports to the Ascend network. In the core of the network, the switches function as IP hardware forwarding engines. The advantages to implementing IP switching technology over traditional routing include lower-layer packet handling, improved traffic management and throughput, increased performance, and end-to-end Quality of Service (QoS).

In existing Internet Service Provider (ISP) networks, the addition of IP switching allows service providers to optimize data traffic flow by eliminating the need for all data packets to flow through the core router. IP switching also eases the management and control duties of the core routers by reducing the number of routing sessions, eliminating IP table lookups, and in some cases, removing the need for the core router completely.

Core routers are still required for connecting LAN-based servers.

Ascend's Implementation of IP Switching

Ascend adds its IP Navigator software to existing multiservice WAN platforms enabling service providers to offer standard or enhanced IP services based on end-to-end Quality of Service (QoS). IP Navigator is a software upgrade to the NMS for Ascend's multiservice switch platforms. (For specific hardware and software requirements, refer to the software release notes that accompany your IP Navigator software package.)

IP Navigator enables a B-STDX and/or CBX switch on the edge of a WAN to run standard IP routing protocols. CBX and/or B-STDX switches are on the edge of the Ascend cloud forwarding packets based on the IP address of the frames. Frame relay, ATM, and Ethernet interfaces are supported. Inside the cloud, a CBX 550 can be used to provide a high-speed ATM backbone, whereby packets are *switched* over automatically established virtual paths.

IP Forwarding

IP forwarding decisions are based on routes obtained via standard routing protocols running on the switch. Inside the cloud, OSPF is used as the Interior Gateway Protocol (IGP). The Border Gateway Protocol (BGP) is used as the Exterior Gateway Protocol (EGP), learning routes to networks in other autonomous systems (AS). You can use RIP, OSPF, or static routing on the links to the CPE Routers to learn what networks are reachable through them.

Routing Protocols

IP Navigator supports a variety of IP routing protocols that are required to communicate with traditional routers. IP Navigator includes all the necessary protocols a service provider needs to offer Internet and Intranet services. These protocols include:

- IP
- OSPF
- RIP
- RIP-2
- BGP-4

Interior Gateway Protocols (IGP)

RIP, RIP-2, and OSPF are interior gateway protocols (IGP). An IGP is used to develop the routing tables within a network that is administered by one company or organization. RIP is still widely used in smaller IP networks.

OSPF is the routing protocol that is typically used in new or large IP networks. An expanded version of OSPF is part of the Virtual Network Navigator (VNN), the connection-oriented routing technology used in Ascend switches.

Exterior Gateway Protocols (EGP)

BGP is an exterior gateway protocol (EGP) that exchanges routing information between autonomous systems (ASs). An AS is a set of routers having a single routing policy running under a single technical administration. BGP advertises routes between external BGP neighbors or peers, unlike Interior Gateway Protocol (IGP), which advertises routes within the same autonomous system, such as over OSPF and RIP. When you configure a list of BGP neighbors and networks, you enable these peers and networks to exchange routing information with the BGP-configured switch. See Figure 8-1 on page 8-2 for an example of AS relationships.

The Internet is a collection of autonomous systems. Interconnections among ASs typically do not use interior gateway protocols, instead they use protocols that are classified as exterior gateway protocols (EGP), such as BGP.

Internet and Transport Protocols

IP Navigator supports the following Internet and transport protocols:

- Address Resolution Protocol (ARP)
- File Transfer Protocol (FTP)
- Internet Control Message Protocol (ICMP)
- Internet Protocol (IP)
- Inverse Address Resolution Protocol (InARP)
- Simple Network Management Protocol (SNMP)
- Telnet Protocol (Telnet)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

Multicast Protocols

Multicasting allows a single packet of information to be sent to multiple destinations. Audio and videoconferencing are natural applications for this type of connection. In a future release, IP Navigator will support Internet Group Multicast Protocol (IGMP), Multicast OSPF (MOSPF), and Distance Vector Multicast Routing Protocol (DVMRP) for multicasting.

Exchanging Route Table Information

In any routed network, routers learn about the topology of the network by exchanging routing information. In an IP Navigator network the same information must be exchanged so that every router-enabled switch shares the same network view. IP Navigator uses Ascend's Virtual Network Navigator (VNN) for this purpose. Essentially, VNN is OSPF with Ascend-proprietary extensions.

Each switch running IP Navigator in the Ascend network communicates with every other switch running IP Navigator. Route tables are maintained in each switch and a master table is maintained in the switch's control processor (CP) for B-STDX 8000/9000 switches and in the switch processor (SP) for CBX 500 switches. Each I/O processor (IOP) module stores routing table information, eliminating a single point of failure in the switch.

Mapping Routes to Virtual Circuits

The process of establishing and mapping routes to virtual circuits is the role of the Virtual Network Navigator (VNN). VNN establishes virtual circuits between entry and exit points in the network. The virtual circuits are then mapped to routes based on the egress node. The process of establishing the virtual circuits is automatic. Topology changes, such as the addition of a new switch, trigger a recalculation of all virtual circuits. More importantly, VNN continually monitors the performance of each virtual circuit, recalculating a new path if a better one exists. The monitoring process is a standard function of VNN and its functions are expanded to include IP Navigator.

Multipoint-to-Point Tunneling

IP Navigator uses Multipoint-to-Point Tunnels (MPTs) to interconnect the routers and switches. An MPT can be thought of as the inverse of the point-to-multipoint virtual circuits that are used to transmit packets from one source to multiple destinations. MPT allows multiple nodes to share the same circuit to transmit to a single destination. Connections of this type are commonly used in multicast applications such as video distribution. The MPT reverses the direction of the information flow. MPTs are a unique feature provided by Ascend.

On adding a new switch to an IP Navigator network, the switch establishes MPTs to all other switches running IP Navigator in the network. This provides every switch with a path to the new switch.

MPTs are established using the best route available through the network. The connections do not have a QoS guarantee. Any information transferred over the link is sent with the lowest level priority on the link. For an ATM link, the information is classified as unspecified bit rate/available bit rate (UBR/ABR). The term *best effort* is used to describe this type of service.

Establishing MPT Circuits

Switches in an IP Navigator network automatically establish MPT circuits upon startup. However, you must configure one MPT parameter. For more information on this, see “Configuring IP Parameters” on page 9-16. VNN calculates the best path from each node to the new switch using the extended version of OSPF. These paths are used to form the MPT. Major network changes cause VNN to recalculate the MPTs. If configured to do so, VNN will continually monitor the MPTs and recalculate them based on performance. VNN treats the MPTs the same way it treats any other circuit connection. No special configuration is required for MPT monitoring.

About MPT Paths

The reverse multipoint tunnel (MPT) enables IP Navigator to switch connectionless protocols (e.g. IP) across the Ascend switch network. MPTs provide an efficient, fault-tolerant, high-performance protocol switching layer that is scalable to 400 switches in a network. MPTs run across direct or optimum trunks which connect CBX 500s and B-STDX 9000s.

The reason for using MPTs over point-to-point tunnels is because MPTs need less circuits than point-to-point tunnels. MPTs require the number of circuits to be equal to the number of nodes, whereas point-to-point tunnels require the number of circuits to be equal to the number of nodes squared.

Establishing End-To-End QoS

The basic service offered by IP Navigator is considered best effort. IP traffic is transmitted with the lowest level of priority. If this best effort service is considered unacceptable, IP Navigator enables you to provision virtual circuits for a specific route. When you configure a circuit for a specific route, you can assign a desired QoS. As with all services offered on Ascend's switch platforms, the QoS is guaranteed on an end-to-end basis.

Configuration and Management

With the addition of IP Navigator, NavisCore and associated network management server products provide the required support for all IP switching features.The protocols required to configure IP switching include: IP, OSPF, RIP, and BGP. In addition IP Navigator adds new monitoring functions to enable network administrators to monitor their IP traffic parameters and routing-table contents.

IP Navigator supports the following standard MIBs:

- MIB II
- OSPF v2 MIB [3]
- BGP-4 MIB [4]
- Routing Table MIB [5]
- RIP v2 MIB [6]

Logical Port Configuration

You can configure the following types of logical ports for IP routing:

- IP logical ports on B-STDX 8000/9000 and CBX 500 switches. IP logical ports are ports that support IP routing.
- IP Server logical ports on the CBX 500. IP Server logical ports provide a method of accepting or transmitting IP traffic to or from a cell-based logical ports.

Table 1-1 lists the logical ports and card types that support IP routing on the B-STDX 8000/9000. Table 1-2 lists the logical ports and card types that support IP routing on the CBX 500.

See Chapter 3, “Configuring IP Logical Ports and IP Servers” for further details about logical port configuration.

Table 1-1. Logical Ports Supporting IP Routing on the B-STDX 8000/9000

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE FR NNI	Frame cards[a]	RFC1490	InARP (RFC 1293) ARP (RFC1490) Static configuration
PPP	Frame cards[a]	PPP	Static configuration
ATM UNI DTE ATM UNI DCE	Frame cards[a]	RFC 1483	InATMARP ATMARP
ATM UNI DTE ATM UNI DCE	ATM cards[b]	RFC 1483	InATMARP ATMARP
Ethernet	2-port Ethernet card[c]	IEEE SNAP Ethernet II	ARP

[a] Frame Cards = UIO, 4-T1, 4-E1, DSX-10, HSSI, Ch T3

[b] ATM Cards = ATM UNI Rev C, ATM CS, ATM DS3, ATM E3, ATM OC3

[c] Please contact your Sales Representative for information regarding the availability of the 2-Port Ethernet module.

Table 1-2. Logical Ports Supporting IP Routing on the CBX 500

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE FR NNI	6-Port DS3 Frame card	RFC 1490	ARP Inverse ARP
Ethernet	4-port Ethernet card	IEEE SNAP Ethernet II	ARP
ATM UNI DTE ATM UNI DCE	ATM cards[b] with an IP Server PVC connection	RFC 1483	ARP Inverse ARP

2

Configuring Ethernet Logical Ports

This chapter describes how to configure logical ports for the following cards:

- 4-port Ethernet for use on the CBX
- 2-port Ethernet for use on the B-STDX

Each of these Ethernet cards support speeds of up to 100 Mbps full duplex.

Please contact your Sales Representative for information regarding the availability of the 2-Port Ethernet module.

Prerequisites

Before you configure an Ethernet logical port, check to make sure that you have:

- Set the Ethernet card's attributes.
- Defined the physical ports on which the logical port(s) will reside.

For more information about these two tasks, refer to the *NavisCore Physical Interface Configuration Guide*.

Accessing the Logical Port Functions

To access the Logical Port functions in NavisCore:

1. Select the switch to which you want to add a logical port.
2. Log in to NavisCore using either a provisioning or operator password.
3. From the Administer menu, select Ascend Parameters ⇒ Set Parameters.

The Switch Back Panel dialog box appears. Figure 2-1 illustrates the Switch Back Panel dialog box for a CBX 500.

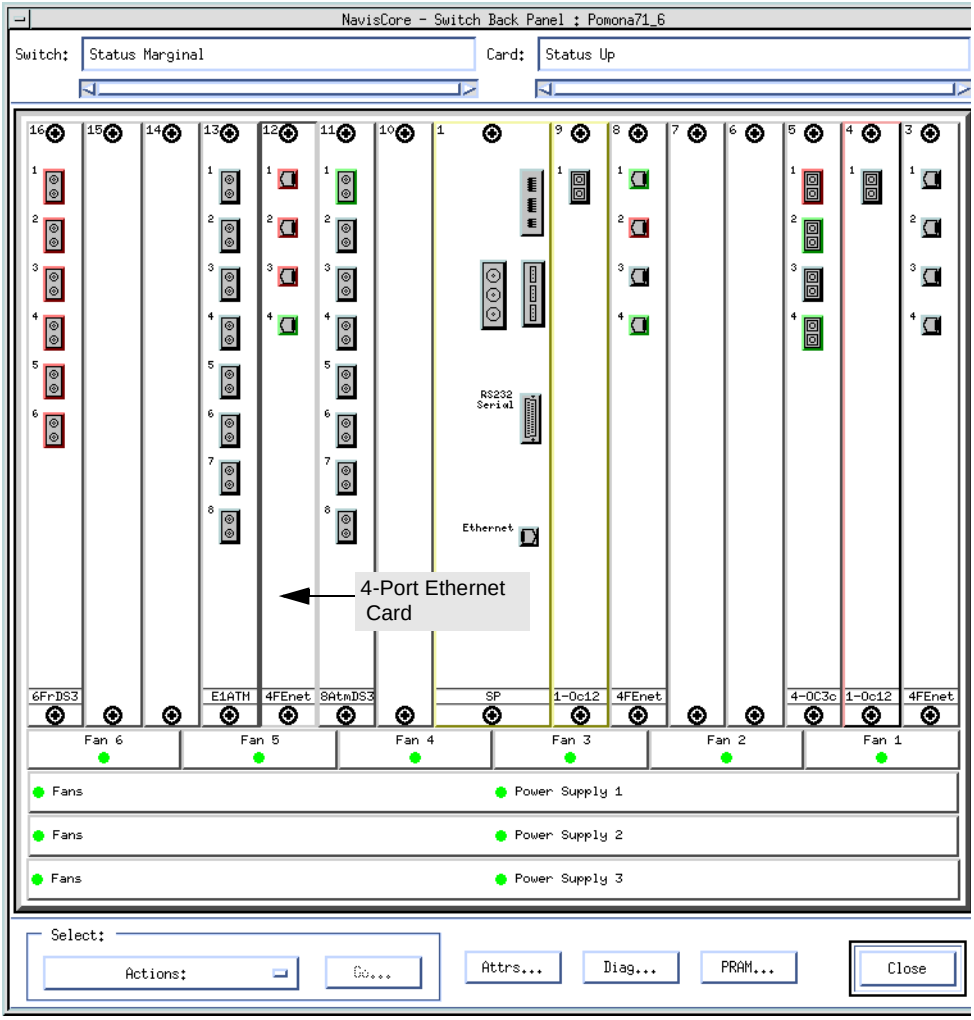

Figure 2-1. Switch Back Panel (CBX) Dialog Box

4. Select the physical port you want to configure. Choose the Attrs... button. The Set Physical Port Attributes dialog box appears (Figure 2-2).

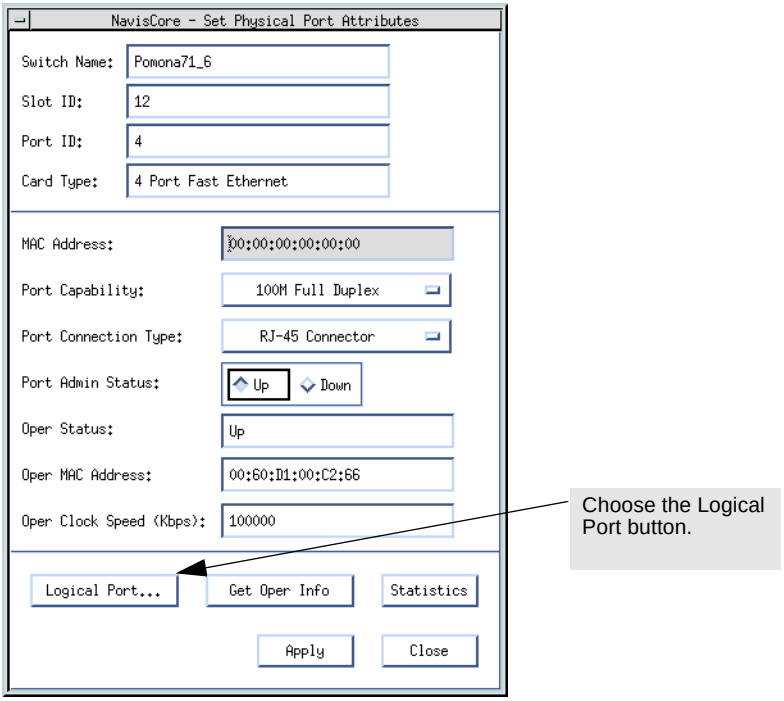

Figure 2-2. Set Physical Port Attributes

5. Choose Logical Port. The Set All Logical Ports in PPort dialog box appears (Figure 2-3).

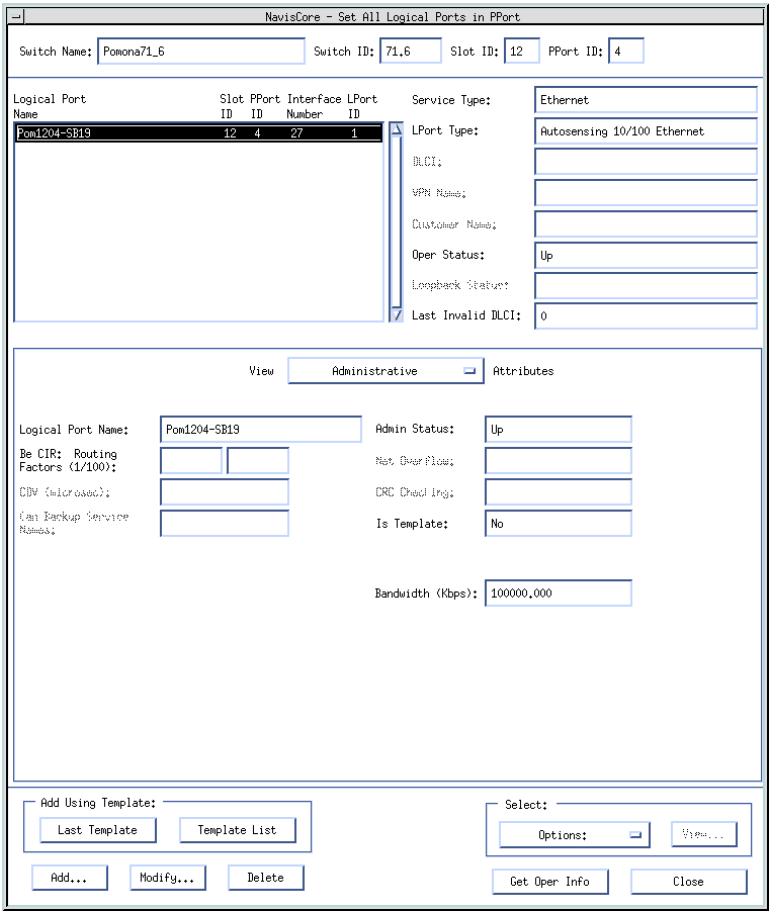

Figure 2-3. Set All Logical Ports in PPort Dialog Box

About the Set All Logical Ports Dialog Box

The Set All Logical Ports In PPort dialog box displays information about an existing logical port or enables you to add a new logical port. It also provides several command buttons that you can use to access additional logical port functions, such as add, modify, and delete logical ports.

Table 2-1 describes the Set All Logical Ports in PPort command buttons.

Table 2-1. Set All Logical Ports in PPort Command Buttons

Field/Command	Action/Description
Add	Adds a new logical port.
Modify	Modifies the selected logical port. The Modify command displays dialog boxes which are similar to those displayed when you Add a logical port; however, you cannot modify the logical port name and the logical port type.
Delete	Deletes the selected logical port.
Get Oper Info	Displays a brief status message of the logical port state.
Add Using Template	If you have already defined a logical port configuration and saved it as a template, use this option to define a new logical port using similar parameters. • Choose Last Template to use the last template you defined for this switch. • Choose Template List to display a list of templates previously defined for this map.
Options	 Use the Select: Options button to select the following logical port options for Ethernet logical ports. **IP Parameters** — Displays the Set IP Parameters dialog box (Figure 3-5 on page 3-8). **Statistics** — Displays the summary statistics for the selected logical port. For more information about summary statistics, see the *NavisCore Diagnostic and Troubleshooting Guide.* **Diagnostics** — Accesses diagnostic tests for the selected logical port. For more information about diagnostics, see the *NavisCore Diagnostic and Troubleshooting Guide.* Once you select an option from this list, choose View to access the information.

6. Choose Add to define a new logical port. The Add Logical Port Type dialog box appears (Figure 2-4).

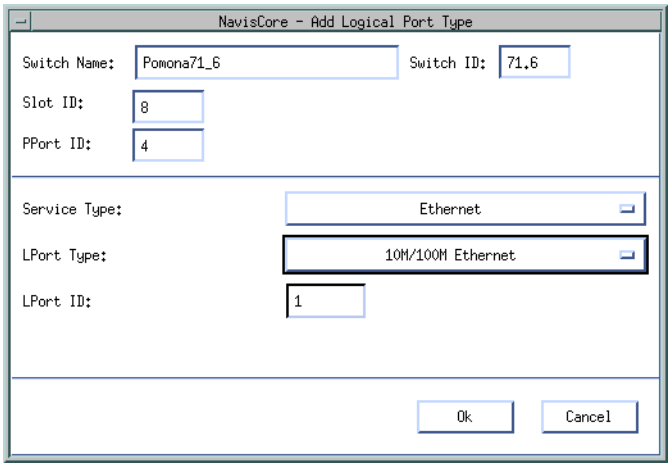

Figure 2-4. Add Logical Port Type Dialog Box

7. Accept the displayed values and choose OK. The Add Logical Port dialog box appears.

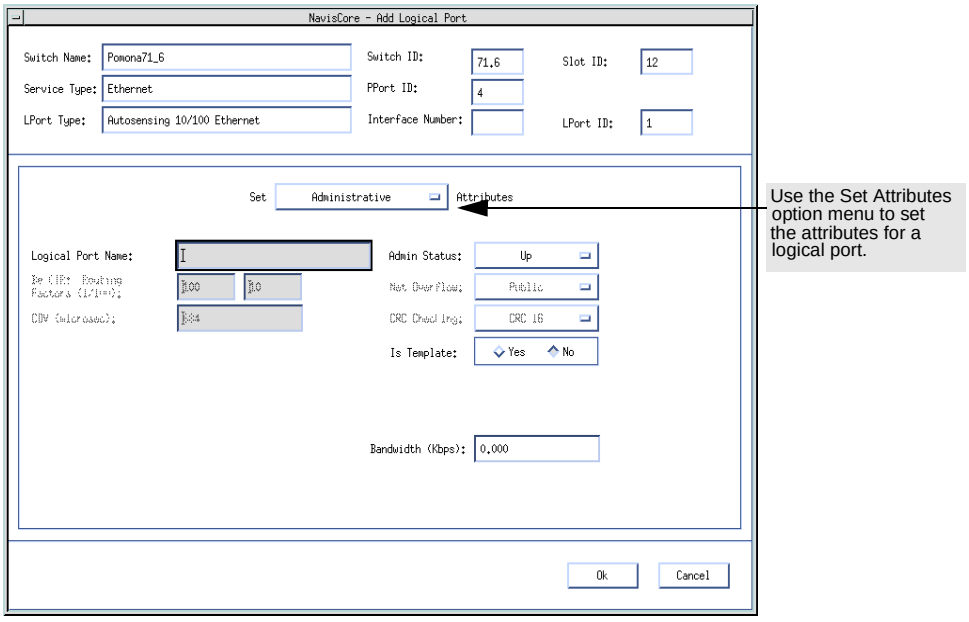

Figure 2-5. Add Logical Port Dialog Box

The Set Attributes Options Menu

When you define a new logical port, the Add Logical Port dialog box displays a Set Attributes option menu that enables you to set different attributes for each type of logical port. Attributes that you can set include:

Administrative — Sets the admin status, net overflow, and bandwidth parameters.

Trap Control — Sets the congestion threshold percentage in which traps are generated and the number of frame errors per minute for each logical port.The supported logical port types are different for each I/O module.

Ethernet Frame — Sets the encapsulation type for transmitted IP frames on an Ethernet port.

Administrative Attributes

Use the Set Administrative Attributes option to complete the fields described in Table 2-2.

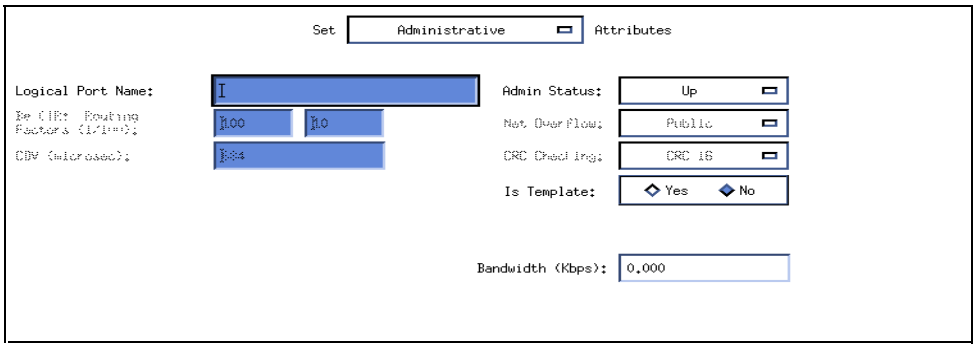

Figure 2-6. Administrative Attributes for Ethernet Logical Ports

Table 2-2. Administrative Attributes (Ethernet Ports) Fields

Field	Action/Description
Logical Port Name	Enter a unique alphanumeric name for this port. NavisCore uses this name to reference the logical port.
Admin Status	Set the Admin Status to Up (the default) to make the port active. Set the Admin Status to Down to make the port inactive.
Is Template	*(Optional)* Saves these settings as a template to configure another logical port with similar options. To create a template, choose Yes.
Bandwidth	The bandwidth for this logical port. The default is 100,000 Kbps, which is the optimal physical clock rate.

Trap Control Attributes

Use the Set Trap Control Attributes option to complete the fields described in Table 2-3

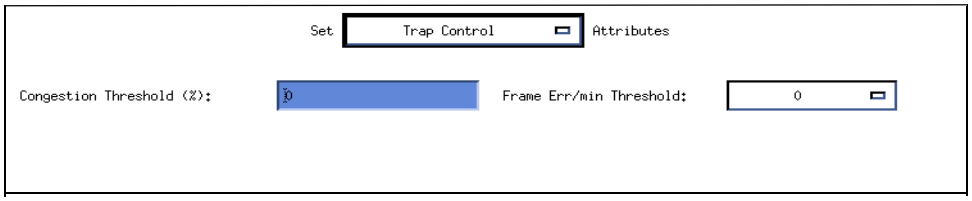

Figure 2-7. Trap Control Attributes for Ethernet Logical Ports

Table 2-3. Set Trap Control Attributes (Ethernet Ports) Fields

Field	Action/Description
Congestion Threshold (%)	Enter a value between 0 and 100 to indicate the threshold percentage for generating and sending traps to the NMS for this logical port. A congestion trap is generated and sent to the NMS if the rate of congestion over a one-minute period exceeds the percentage value you enter. Adjust the entered value according to how sensitive this port needs to be to network congestion. Options include: *Low* – Generates a trap at the first sign of congestion. *High* – Generates traps for serious network congestion. *Zero (default)* – Disables the congestion threshold. If you enter zero, no traps are generated for this logical port.
Frame Err/Min Threshold	Enter a value from 0 to 16384 to configure the threshold of frame errors on this logical port. If the number of frame errors received in one minute exceeds the specified number, a trap is sent to the NMS. Adjust this value according to how sensitive this port needs to be to frame errors. Options include: *Low* – Port is sensitive to frame errors. *High* – Generates traps when a significant number of frame errors occurs within a one-minute period. *Zero (default)* – Disables this feature, which prevents traps from being generated for this logical port.

Ethernet Attributes

Use the Set Ethernet Frame Attributes option to complete the fields described in Table 2-4.

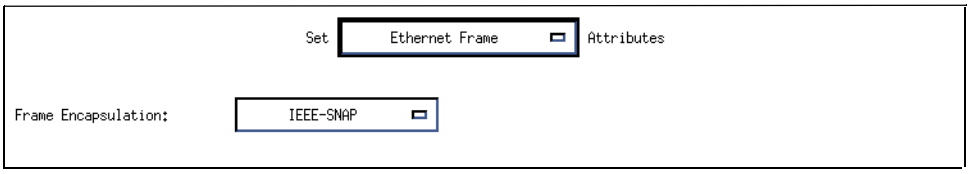

Figure 2-8. Ethernet Attributes for Ethernet Logical Ports

Table 2-4. Set Ethernet Frame Attributes (Ethernet Ports) Fields

Field	Action/Description
Encapsulation Type	Select the Ethernet frame encapsulation type: *Ethernet-II – (Default)* The frame type used in Novell NetWare networks. The following figure illustrates an Ethernet II frame. The value of the Ethernet Type (specified in bytes 13 and 14) indicates the frame type for the packet. If the value of Ethernet type is greater than or equal to hexadecimal value 0x0600 (decimal value 1800), the packet uses an Ethernet II frame type. 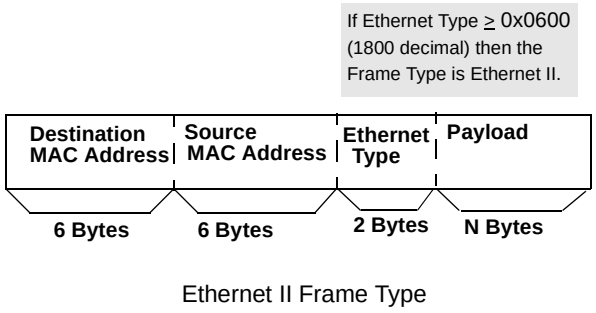 Ethernet II Frame Type *IEEE-SNAP* – An IEEE standard frame type that is 8 bytes larger than the Ethernet II type.The following figure illustrates the IEEE-SNAP Frame. The value of the Length (specified in bytes 13 and 14) indicates the frame type for the packet. If the value of the Length is less than or equal to hexadecimal value 0x05DC (decimal value 1500), the packet uses an IEEE-SNAP frame type. 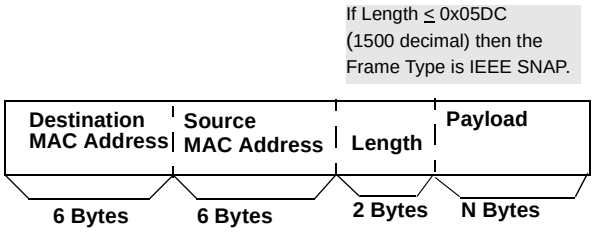 IEEE SNAP Frame Type

3

Configuring IP Logical Ports and IP Servers

This chapter describes how to configure:

- **IP logical ports** on B-STDX 8000/9000 and CBX 500 switches. IP logical ports are ports that support IP routing. See Figure 3-1 on page 3-4 for a summary of the IP logical port configuration process.
- **IP server logical ports** on the CBX 500 switch. IP server logical ports provide a method of accepting or transmitting IP traffic on a cell-based card. See Figure 3-13 on page 3-20 for a summary of the IP server logical port configuration process.

Table 3-1 lists the logical ports that support IP routing.

Table 3-1. Logical Ports that Support IP Routing

Switch	Logical Port	Card Types
B-STDX 8000/9000	FR UNI-DCE FR UNI-DTE FR NNI	Frame cards[a]
B-STDX 8000/9000	PPP	Frame cards[a]
B-STDX 8000/9000	ATM UNI DTE ATM UNI DCE	Frame cards[a]
B-STDX 8000/9000	ATM UNI DTE ATM UNI DCE	ATM cards[b]
B-STDX 8000/9000	Ethernet	2-port Ethernet cards[c]
CBX 500	Frame Relay	6-Port DS3 Frame card
CBX 500	Ethernet	4-port Ethernet card
CBX 500	ATM UNI DTE ATM UNI DCE	ATM cards[b] with an IP Server PVC connection

[a] Frame Cards = UIO, 4-T1, 4-E1, DSX-10, HSSI, Ch T3

[b] ATM Cards = ATM UNI Rev C, ATM CS, ATM DS3, ATM E3, ATM OC3

[c] Please contact your Sales Representative for information regarding the availability of the 2-Port Ethernet module.

Prerequisites

Prior to configuring IP services, verify that the following tasks are complete.

- Create a network map.
- Configure the switch parameters. If you are configuring an IP logical port on a CBX 500, you must install and configure a 6-port DS3 Frame card or a 4-port Ethernet card.
- Configure the physical port parameters.
- If you are configuring an IP logical port on a B-STDX 8000/9000, configure the logical port for Frame Relay, ATM, or Ethernet service.

For more details about these tasks, see Chapter 2, "Configuring Ethernet Logical Ports" and the following guides:

- *NavisCore Frame Relay Configuration Guide* for information about configuring logical ports for frame relay service.
- *NavisCore ATM Configuration Guide* for information about configuring logical ports for ATM service.
- *NavisCore Physical Interface Configuration Guide* for information about configuring cards and physical ports.

About IP Addresses

When you specify the IP address, you must specify the type of IP forwarding the logical port will use. The following two types of IP forwarding are automatically enabled by default:

Unicast — Enables IP forwarding from this logical port to a unicast address.

Broadcast — Enables IP forwarding from this logical port to a broadcast address.

Address Resolution Protocol

A node requires the following information to communicate with another node:

- IP address of the destination node
- Hardware address of the destination node (DLCI for Frame Relay and VPI/VCI for ATM)

IP services uses one of the following protocols to resolve an unknown hardware or IP address:

Address Resolution Protocol (ARP) — Is used for Ethernet configurations, when an IP address of a given destination is known, but the destination hardware address (DLCI or VPI/VCI) is not.

Inverse Address Resolution Protocol (InARP) — Is used for Frame Relay and ATM configurations, when the destination hardware address (DLCI or VPI/VCI) is known, but the destination IP address is not.

The ARP table resides in the CP/SP memory. An ARP entry is stored for 25 minutes (the same amount of time as a BSD IP stack). All statically configured ARP entries are stored in PRAM. If there is a change in the ARP table, it is sent to the IOP's.

Configuring IP Logical Ports

Figure 3-1 illustrates the steps for configuring an IP logical port on a B-STDX 8000/9000 or on a CBX 500.

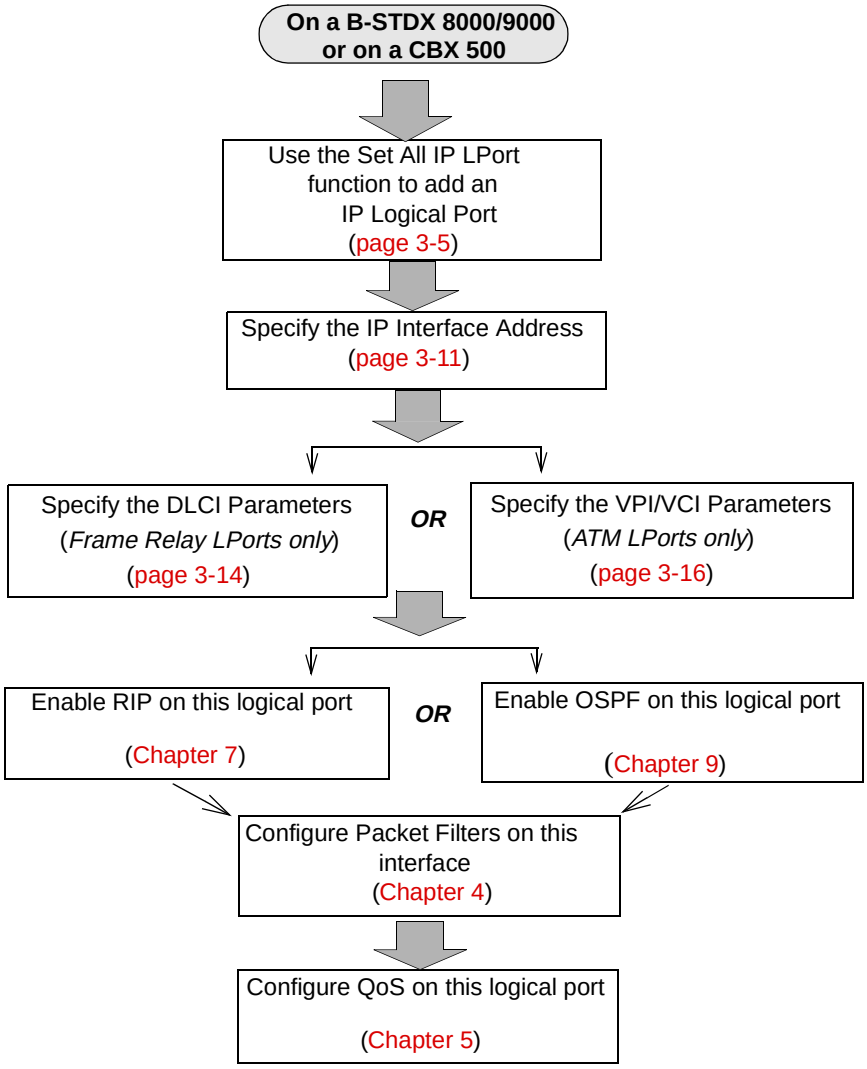

Figure 3-1. IP Logical Port Configuration Process

Accessing Logical Port Parameters

The following section describes how to access the screens that you will use to configure the IP Logical Port Parameters. You can use either of the following methods to access the Set IP Parameters dialog box:

- From the NavisCore Menu. See the following section, "Accessing the Set IP Parameters Dialog Box from the NavisCore Menu" for details about this method of access.
- From the Set All Logical Ports in PPort dialog box. See "Accessing the Set IP Parameters Dialog Box from the Set All Logical Ports Dialog Box" on page 3-6.

Accessing the Set IP Parameters Dialog Box from the NavisCore Menu

To access the Set IP Parameters dialog box from the NavisCore menu:

1. Select the appropriate switch icon from the network map.
2. From the Administer menu, select Ascend IP Parameters ⇒Set All IP Lports. The Set All IP LPorts dialog box appears (Figure 3-2).

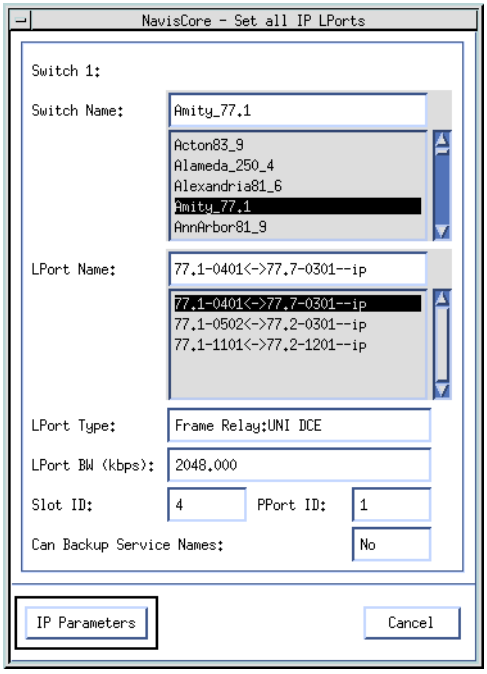

Figure 3-2. Set All IP LPorts Dialog Box

3. Select the LPort name from the list of LPorts.
4. Choose IP Parameters.

If no IP interfaces have already been defined for this switch, the Set IP Parameters dialog box appears (Figure 3-3).

If IP interfaces were previously defined for this switch, the Set IP Parameters dialog box shown in Figure 3-5 appears.

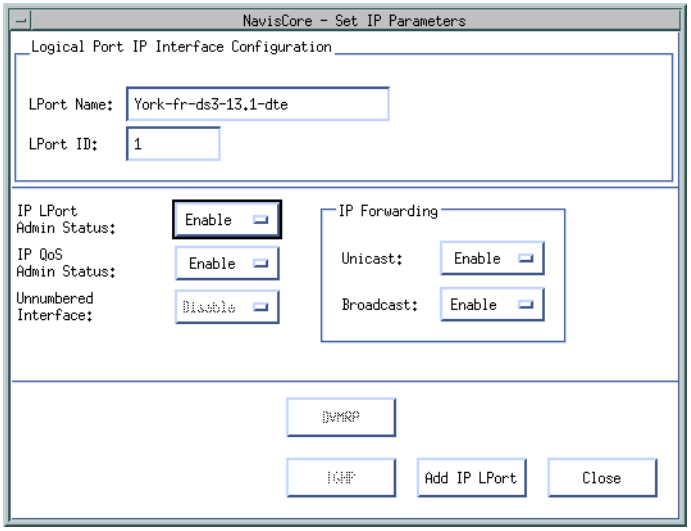

Figure 3-3. Set IP Parameters Dialog Box

See "Adding an IP Logical Port" on page 3-8 for instructions on adding an IP LPort from the Set IP Parameters dialog box.

Accessing the Set IP Parameters Dialog Box from the Set All Logical Ports Dialog Box

To access the Set IP Parameters dialog box from the Set All Logical Ports in PPort dialog box:

1. From the network map select the appropriate switch icon.
2. From the Administer menu select Ascend Parameters ⇒Set Parameters. The Switch Back Panel appears.
3. Select the physical port and choose Attrs... The Set Physical Port Attributes dialog box appears.
4. Choose Logical Port. The Set All Logical Ports in PPort dialog box appears (Figure 3-4).

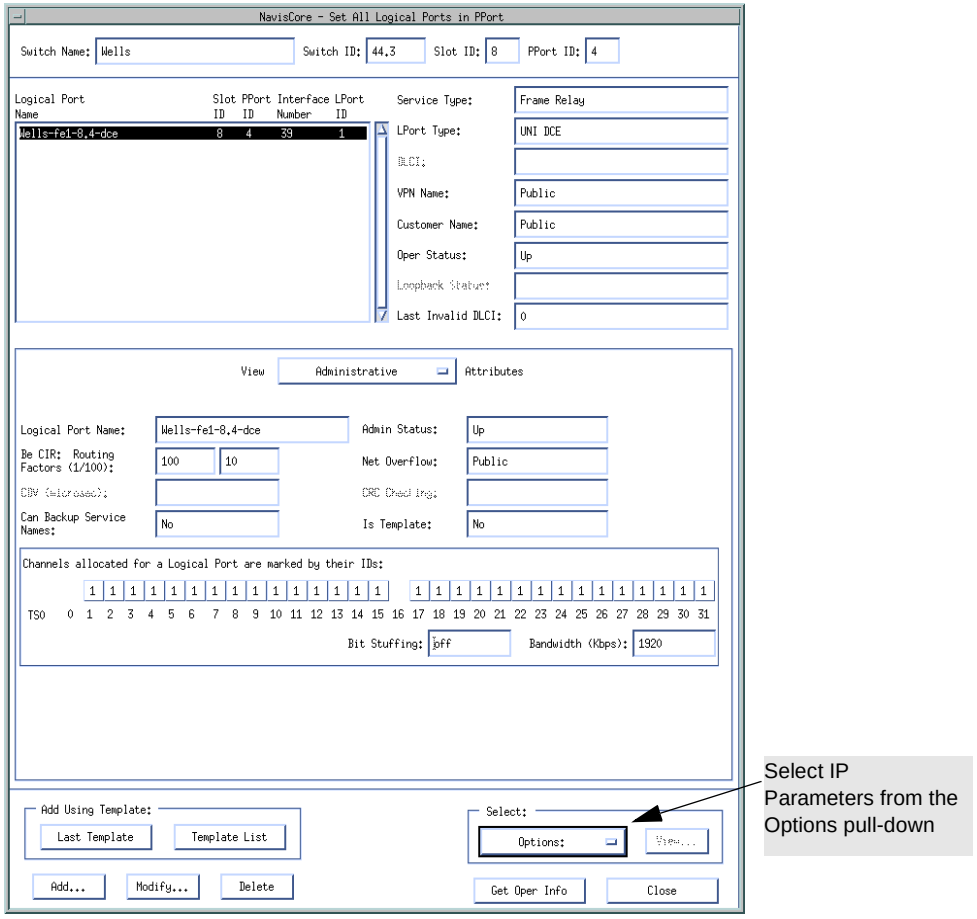

Figure 3-4. Set All Logical Ports in PPort

5. Select the logical port name from the list.
6. Select IP Parameters from the Options pull-down menu.
7. Choose Set. The Set IP Parameters dialog box appears (Figure 3-3 on page 3-6).

Adding an IP Logical Port

To add an IP logical port:

1. Choose Add IP LPort from the Set IP Parameters dialog box (Figure 3-3 on page 3-6). The second Set IP Parameters dialog box appears (Figure 3-5).

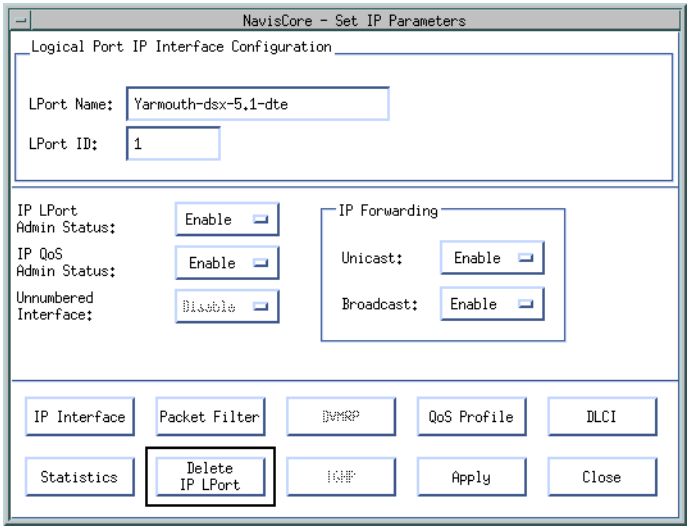

Figure 3-5. Second Set IP Parameters Dialog Box

See Table 3-2 for a description of each of the buttons on the Set IP Parameters dialog box.

Table 3-2. Set IP Parameters Buttons

Button	Function
IP Interface	Displays the Set IP Interface Addresses dialog box enabling you to configure the IP interface address. See page 3-11 for details.
Packet Filter	Displays the Assign Logical Port IP Filter dialog box enabling you to specify inbound and outbound packet filters. See "Configuring IP Packet Filters" on page 4-2 for more details on this function.
QoS Profile	Displays the Associate LPort QoS Profile dialog box enabling you to add and associate Quality of Service profiles. See Chapter 5, "Provisioning IP Quality of Service" for more details on this function.
DLCI	*(For Frame Relay modules only)* Displays the IP Protocol Connection ID dialog box enabling you to specify the Data Link Connection Identifier (DLCI) for the IP logical port. See "Setting the DLCI for Frame Relay Logical Ports" on page 3-14 for more details on this function.
VPI/VCI	*(For ATM modules only)* Displays the IP Protocol Connection ID dialog box enabling you to specify the Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) for the IP logical port. See"Setting the VPI/VCI for ATM Logical Ports" on page 3-16 for more details on this function.
Statistics	Displays the IP Lport Statistics dialog box. See the *NavisCore Diagnostic and Troubleshooting Guide* for more information about IP logical port statistics.
Delete IP Lport	Choose this option to delete the IP configuration values for this logical port so that the port is no longer an IP logical port.
Apply	Applies any modifications made to the IP logical port parameters. If you make changes to the IP logical port parameters on the Set IP Parameters dialog box, the changes are not actually made until you choose Apply.

2. Specify the necessary IP Parameter values listed in Table 3-3.

Table 3-3. IP Parameter Fields

Field	Action/Description
Lport Name	Displays the name assigned to the LPort at configuration. If you plan to use this logical port as a QoS PVC, it is suggested that the Lport Name identify the port as a QoS logical port. When you later use this logical port to associate to the QoS PVC, you will have to select the logical port from a list of Lport Names. Chapter 5, "Provisioning IP Quality of Service" provides details about QoS.
Lport ID	Displays the ID number that uniquely identifies each logical port.
IP LPort Admin Status	Select one of the following options: *Enable* – Indicates that the port is activated for IP services. *Disable* – Indicates that the port has never been activated for IP services or that the port is offline for diagnostics. A logical port card with an IP LPort Admin Status of *Disable* is not operational for IP routing.
IP QoS Admin Status	Select one of the following options: *Enable* – Enables the use of QoS flow profiles for the logical port. *Disable* – Disables the use of QoS flow profiles.
Unnumbered Interface	Select one of the following options: *Enable* – Indicates that this IP logical port is not part of a subnet. It does not have a specific address and instead uses the router ID as its source address in IP packets that originate from the interface and are forwarded out of the interface. (The router ID is always the internal address, regardless of whether or not loopbacks are configured.) *Disable* – Indicates that this IP logical port is part of a subnet.
Unicast	Select one of the following options: *Enable* – Specifies that IP forwarding will be allowed from this logical port to a unicast address. *Disable* – Indicates that IP forwarding will not be allowed from this logical port to a unicast address. The specific unicast addresses are specified for each IP interface. See the Unicast Address descriptions in Table 3-5 on page 3-13 for details on how to specify a unicast address for an IP interface.
Broadcast	Select one of the following options: *Enable* – Specifies that IP forwarding will be allowed from this logical port to a broadcast address. *Disable* – Specifies that IP forwarding is not allowed from this logical port to a broadcast address. The specific broadcast addresses are specified for each IP interface. See the Broadcast Address descriptions in Table 3-5 on page 3-13 for details on how to set a broadcast address for an IP interface.

The next step is to specify the IP interface address for the IP logical port. See the following section, "Setting the IP Interface Address," for details.

Setting the IP Interface Address

To specify the IP Interface Address:

1. From the Set IP Parameters dialog box (Figure 3-3 on page 3-6) choose IP Interface. The Set IP Interface Addresses dialog box appears (Figure 3-6).

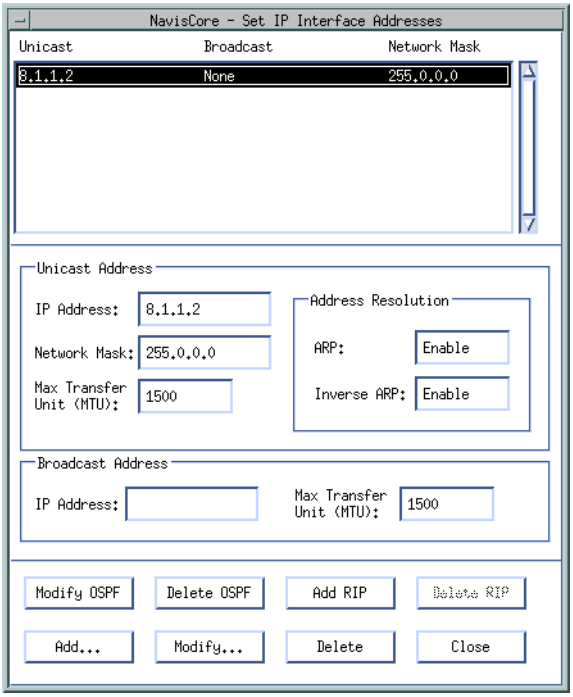

Figure 3-6. Set IP Interface Addresses Dialog Box

Table 3-4. Set IP Interface Addresses Buttons

Button	Function
Add OSPF	Displays the Add OSPF Interface dialog box enabling you to specify the OSPF parameters for the logical port. This button appears only if you have not yet specified any OSPF parameters for the logical port.
Add RIP	Displays the Add RIP Interface dialog box enabling you to specify the RIP parameters for the logical port. This button does not appear if you have already configured RIP parameters for the logical port.
Modify OSPF	Displays the Modify OSPF Interface dialog box enabling you to modify the OSPF parameters for the logical port. This button appears only if you have already specified the OSPF parameters for the logical port.
Modify RIP	Displays the Modify RIP Interface dialog box enabling you to modify the RIP parameters for the logical port. This button appears only if you have already specified the RIP parameters for the logical port.

Table 3-4. Set IP Interface Addresses Buttons (Continued)

Button	Function
Delete OSPF	Displays the Delete OSPF Interface dialog box enabling you to delete the OSPF parameters for the logical port. This button appears only if you have already specified the OSPF parameters for the logical port.
Delete RIP	Displays the Delete RIP Interface dialog box enabling you to delete the RIP parameters for the logical port. This button appears only if you have already specified the RIP parameters for the logical port.
Add	Displays the Add Interface Address dialog box.
Modify	Displays the Modify Interface Address dialog box.
Delete	Displays the Delete Interface Address dialog box.

2. Choose Add to add an IP interface address. The Set IP Interface Address dialog box appears (Figure 3-7).

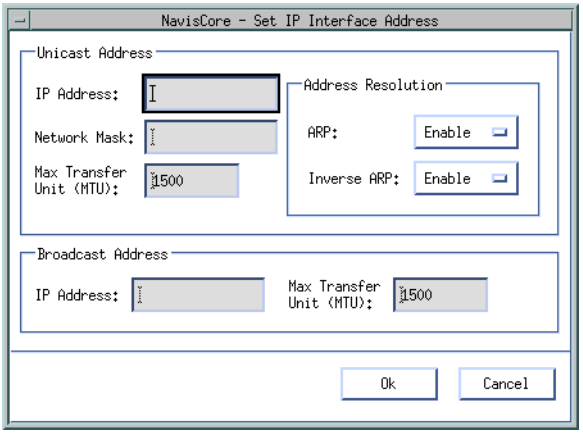

Figure 3-7. Set IP Interface Address Dialog Box

3. Specify the IP Interface Address values described in Table 3-5.

Table 3-5. IP Interface Address Fields

Field	Action/Description
Unicast Address	
IP Address	The IP address for this interface. A maximum of 200 IP interface addresses can be configured on each card. Interface addresses can be distributed across IP logical ports as required.
Network Mask	The mask used to determine the subnet of this IP interface. Once this value is set, you cannot use the Modify Interface Address function to modify the network mask value. In order to change the network mask, you must delete the IP interface and then add a new one using the correct network mask.
Max Transfer Unit (MTU)	The maximum size of a packet that can be sent through the physical port. The default value for this field varies depending on the logical port type as follows: **LPort Type** **Default** ATM 9180 Frame Relay 4096 Ethernet 1500
Address Resolution	
ARP	*(Frame Relay Only)* Select one of the following options: *Enable* – Enables the Address Resolution Protocol (ARP). *Disable* – Disables the ARP. See "Address Resolution Protocol" on page 3-3 for details.
Inverse ARP	Select one of the following options: *Enable* –Enables the Inverse Address Resolution Protocol (InARP). *Disable* – Disables the InARP. See "Address Resolution Protocol" on page 3-3 for details.
Broadcast Address	
IP Address	The address used by this interface for subnet broadcasting.
Max Transfer Unit (MTU)	The maximum size of a packet that can be sent through the physical port. he default value for this field varies depending on the logical port type as follows: **LPort Type** **Default** ATM 9180 Frame Relay 4096 Ethernet 1500

4. Choose OK.

After you assign the IP interface address you can then specify the DLCI (for Frame Relay logical ports) or the VPI/VCI (for ATM logical ports). See the following sections for more information on these tasks:

- "Setting the DLCI for Frame Relay Logical Ports" on page 3-14
- "Setting the VPI/VCI for ATM Logical Ports" on page 3-16

Setting the DLCI for Frame Relay Logical Ports

A data link connection identifier (DLCI) number is a 10-bit address that identifies PVCs. The range for an IP DLCI number is a value from 16 to 991.

To specify the DLCI for Frame Relay Logical Ports:

1. From the Administer menu choose Ascend IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.
2. Select the LPort and choose IP Parameters. The Set IP Parameters dialog box appears.
3. Choose DLCI. The IP Protocol Connection ID dialog box appears (Figure 3-8).

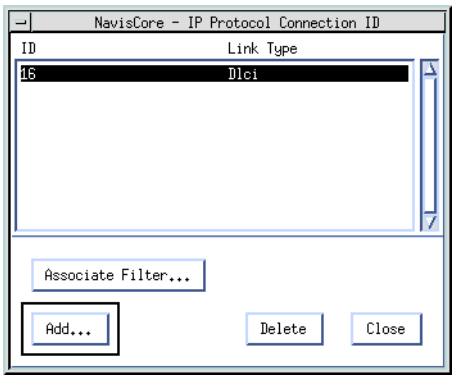

Figure 3-8. IP Protocol Connection ID Dialog Box

The IP Protocol Connection ID dialog box provides the following option buttons:

Button	Function
Add	Displays the Set IP Protocol Connection ID dialog box to enable you to add a DLCI number.
Delete	Deletes a selected DLCI number.
Associate Filter	Displays the Associate IP Circuit Filter List dialog box to enable you to associate a filter with a specific DLCI address. See "Assigning IP Packet Filters to Circuits" on page 4-18 for more information.

4. Choose Add. The Set IP Protocol Connection ID dialog box appears (Figure 3-9).

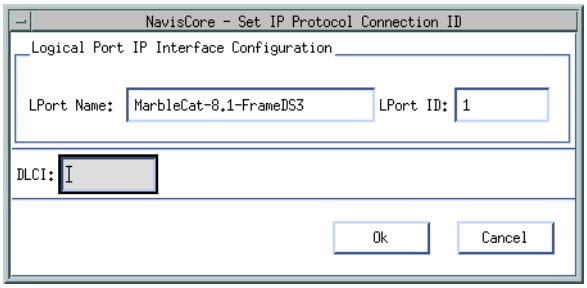

Figure 3-9. Set IP Protocol Connection ID Dialog Box

Specify the field values as described in Table 3-6.

Table 3-6. IP Protocol Connection ID For Frame Relay LPorts Fields

Field	Action/Description
Lport Name	Displays the name assigned to the LPort at the time of configuration. The LPort ID uniquely identifies each logical port within the physical port.
Lport ID	Displays the ID number that uniquely identifies each logical port.
DLCI	The DLCI value for this IP interface. The range for an IP DLCI number is a value from 16-991: 0-15 – is reserved 16-991 – is available for most link management types. If link management is LMI-1, the maximum value is 1007.

Setting the VPI/VCI for ATM Logical Ports

Virtual path identifiers (VPIs) and virtual channel identifiers (VCIs) are addressing identifiers (similar to Frame Relay's DLCI) that route cell traffic. The ATM cell header contains both a VCI and a VPI, which provides an ATM cell with a unique VCI and associates it with a particular virtual path. Every ATM cell uses these VPI/VCI identifiers.

> The VPI and VCI are used only for establishing connections between two ATM entities, not the end-to-end connection.

To specify the VPI and VCI for ATM logical ports:

1. From the Administer menu choose Ascend IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.
2. Select the LPort and choose IP Parameters. The Set IP Parameters dialog box appears.
3. Choose VPI/VCI. The IP Protocol Connection ID dialog box appears (Figure 3-10).

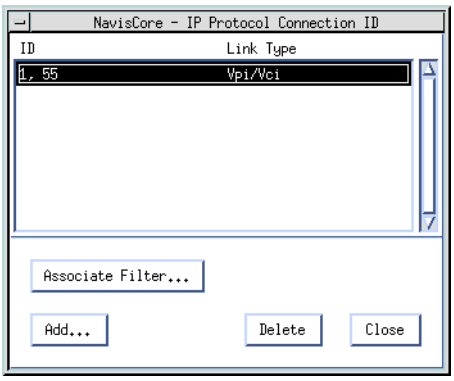

Figure 3-10. IP Protocol Connection ID Dialog Box (ATM LPorts)

The IP Protocol Connection ID dialog box provides the following option buttons:

Button	Function
Add	Displays the Set IP Protocol Connection ID dialog box to enable you to add a VPI/VCI number.
Delete	Deletes a selected VPI/VCI number.
Associate Filter	Displays the Associate IP Circuit Filter List dialog box to enable you to associate a filter with a specific VPI/VCI address. See "Assigning IP Packet Filters to Circuits" on page 4-18 for more information.

4. Choose Add. The Set IP Protocol Connection ID dialog box appears (Figure 3-11).

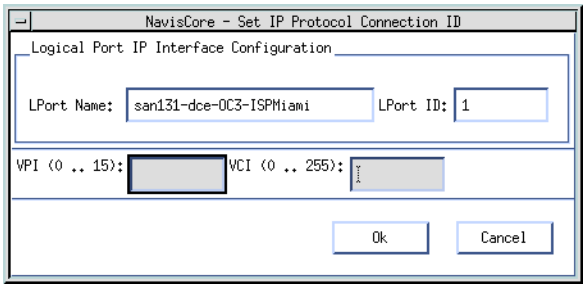

Figure 3-11. Set IP Protocol Connection ID Dialog Box (ATM LPorts)

5. Specify the field values as described in Table 3-7.

Table 3-7. IP Protocol Connection ID For ATM LPorts Fields

Field	Action/Description
LPort Name	Displays the name assigned to the LPort at configuration. The LPort ID identifies the selected logical port.
LPort ID	Displays the unique ID number assigned to the selected logical port.
VPI	The virtual path identifier (VPI). A virtual path (VP) is a group of virtual channels (VCs) carried between two points. VPs provide a way to bundle traffic headed in the same direction. The VPI is an addressing identifier that routes cell traffic. Switching equipment checks the VPI portion of the header to route traffic over certain trunks. This field displays a range of valid values based on the number of valid bits that are configured for this logical port. If the number of valid bits is set to 4, the valid range for the VPI value can be from 0 to 15. If VPI/VCIs already exist on the selected logical port, you can change the number of valid bits, however, the new values must be large enough to support the largest PVC configured for this ATM port. See the *NavisCore ATM Configuration Guide* for a complete description of the valid values for VPI. See"Two Forwarding Engines on Each IP Server Card" on page 3-18 for specific information about VPI values for IP server logical ports.
VCI	The virtual channel identifier (VCI). A virtual channel (VC) is a connection between two communicating ATM entities. This field displays a range of valid values based on the number of valid bits that are configured for this logical port. If the number of valid bits is set to 8, the VCI value can be from 32 to 255 (VCI 0 - 31 are reserved and cannot be used per ATM Forum standards). If VPI/VCIs already exist on the selected logical port, you can change the number of valid bits, however, the new values must be large enough to support the largest PVC configured for this ATM port. See the *NavisCore ATM Configuration Guide* for a complete description of the valid values for VPI. See"Two Forwarding Engines on Each IP Server Card" on page 3-18 for specific information about VCI values for IP server logical ports.

About IP Server Logical Ports on the CBX 500

You can configure logical ports as IP server logical ports on either of the following CBX 500 cards:

- 6-port DS3 Frame Card
- 4-port Ethernet Card

The purpose of an IP server logical port is to provide a method of accepting or transmitting IP traffic from or to a cell-based port. All IP traffic entering and exiting a CBX 500 ATM cell card must be transmitted through an IP server logical port that is configured on either a 6-port DS3 Frame card or on a 4-port Ethernet card.

You can configure up to 14 IP server cards in one CBX 500. The number of IP server cards is only limited by the number of slots in the switch.

Two Forwarding Engines on Each IP Server Card

There are two forwarding engines (FEs) on an IP server card. FEs reassemble cells and perform IP lookups. The NMS identifies the two FEs on an IP server card as Server 1 and Server 2.

Each of the two FEs have hard-coded values for the VPI and VCI bit parameters. The value for the VPI bits parameter is permanently set to 6. For this reason, **the maximum number of IP logical ports that you can create for each IP server card is 64 (0-63).** Therefore, if you define 64 IP logical ports on the first FE of the card (which is identified in the NMS as Server 1), you cannot define any IP logical ports on the second FE (Server 2). You can create any number of IP logical ports (up to a maximum of 64) between the two FEs.

The value for the VCI bits parameter is permanently set to 8. For this reason, you can create up to 225 (2^8 - 31) PVCs for each IP logical port.

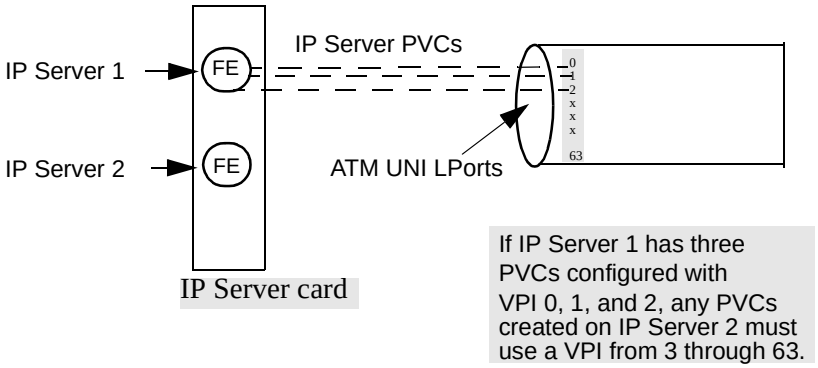

Figure 3-12. VPI Parameters for IP Server Cards

IP Server Logical Ports

IP server logical ports on IP server cards are virtual ports. For this reason, physical port numbers on 6-port DS3 Frame cards start at 7 and physical port numbers on 4-port Ethernet cards start at 5.

IP server cards do not support multiple IP interfaces on the same IP server logical port. For this reason you must create multiple logical ports. You use the Set IP Server function on the Administer menu to configure IP logical ports on a CBX 500. See "Creating an IP Server Logical Port" on page 3-21 for more information.

Bandwidth Allocation

Multiple PVCs can be defined for an IP server logical port.

Logical port bandwidth on an IP server logical port must be sufficient to support all of the PVCs traversing the port. For this reason, before you configure IP server logical ports, **you must plan for the total amount of bandwidth that will be required for all of the PVCs that are associated with the IP server logical port.** If you assign all of the bandwidth to the first IP server logical port, there will be no bandwidth available for PVCs that are configured for subsequent IP Server logical ports.

If you want to create multiple IP server LPorts on the 6-port DS3 Frame card or on the 4-port Ethernet card, you must be aware of the following requirements:

- All of the logical port bandwidth cannot be allocated to the first IP server LPort that you define (a Direct UNI-DCE LPort).
- For subsequent IP logical ports, a VPI start/stop range must be defined. There cannot be any overlap between logical ports. For example, if the first virtual IP server logical port is configured with a VPI start/stop range of 1 and 4, the second virtual IP server logical port cannot use a VPI start/stop range that includes any VPIs already defined.

Configuring IP Server Logical Ports

Figure 3-13 illustrates the steps for configuring an IP server logical port on a CBX 500.

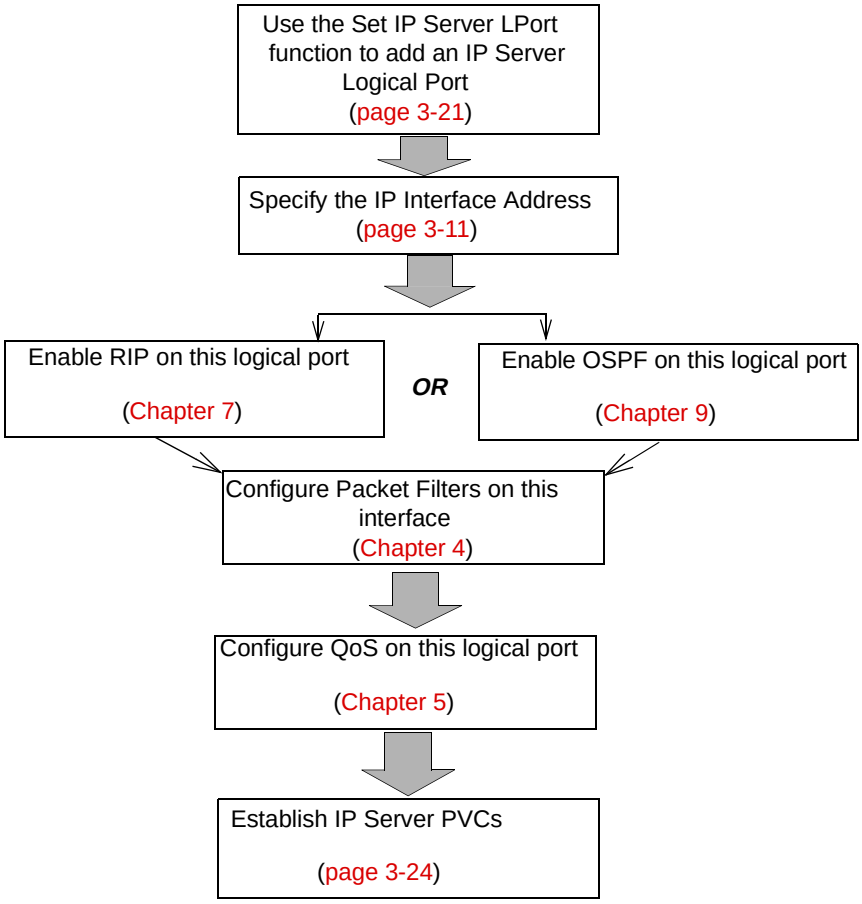

Figure 3-13. Configuring IP Server Logical Ports on the CBX 500

Creating an IP Server Logical Port

To set an IP server logical port from the NavisCore menu:

1. Select the appropriate CBX 500 switch icon from the network map.
2. Select Ascend IP Parameters ⇒ Set IP Servers ⇒ Set IP Server LPorts from the Administer menu. The Show IP Servers dialog box appears (Figure 3-2).

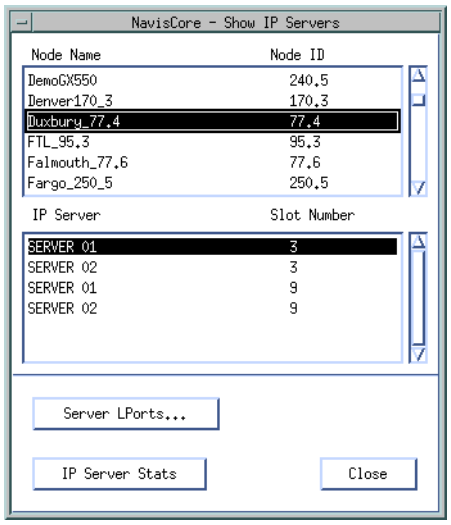

Figure 3-14. Show IP Servers Dialog Box

The Show IP Servers dialog box lists all of the CBX 500 switches in your network. If a switch has one or more IP server cards installed, the cards are listed in the IP Server group box. The switch must have an IP server card installed before you can set an IP logical port. For detailed instructions about how to configure an IP Server card and physical port, see the *NavisCore Physical Interface Configuration Guide*.

The following list describes the buttons on the Show IP Servers dialog box.

Button	Function
Server LPorts	Displays the Set All Logical Ports in IP Server PPort dialog box (Figure 3-15) to enable you to add an IP server logical port.
IP Server Stats	Displays the Physical Port Summary Statistics dialog box to enable you to display the statistics for a selected IP Server physical port. See the *NavisCore Diagnostic and Troubleshooting Guide* for more details about physical port statistics.

3. Select the IP Server name from the list. There are two IP Server FEs available for each IP server card on the switch.
4. Choose Server LPorts.The Set All Logical Ports in IP Server PPort dialog box appears.

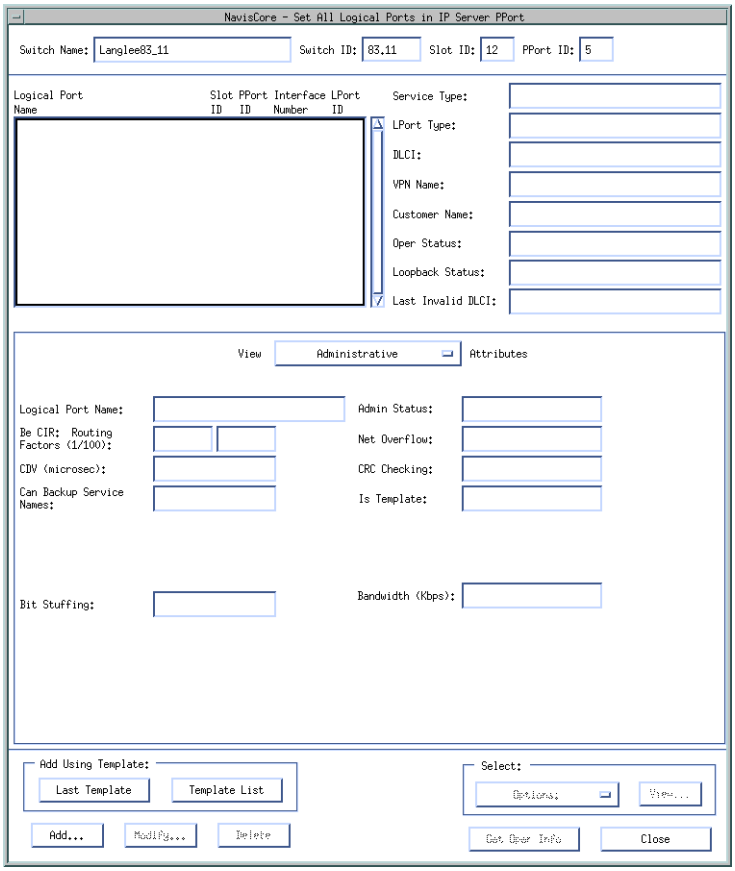

Figure 3-15. Set All Logical Ports in IP Server PPort

5. Choose Add. The Add Logical Port Type dialog box appears.

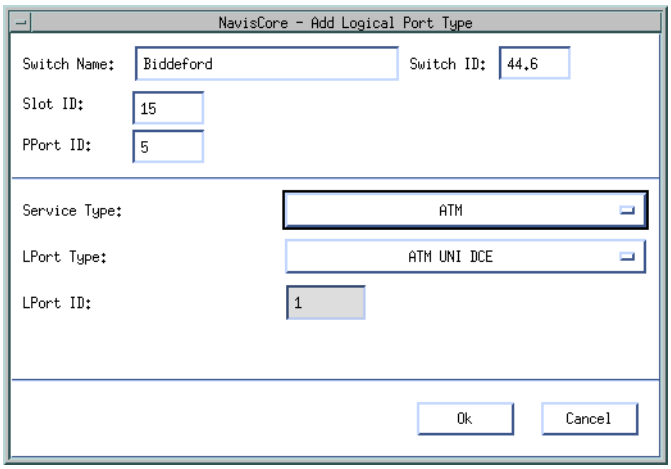

Figure 3-16. Add Logical Port Type

6. Complete the Add Logical Port Type dialog box as follows:

LPort Type — Select ATM UNI DCE.

LPort ID — Defaults to 1 for this type of configuration and cannot be changed.

7. Choose OK. The Add Logical Port dialog box appears.

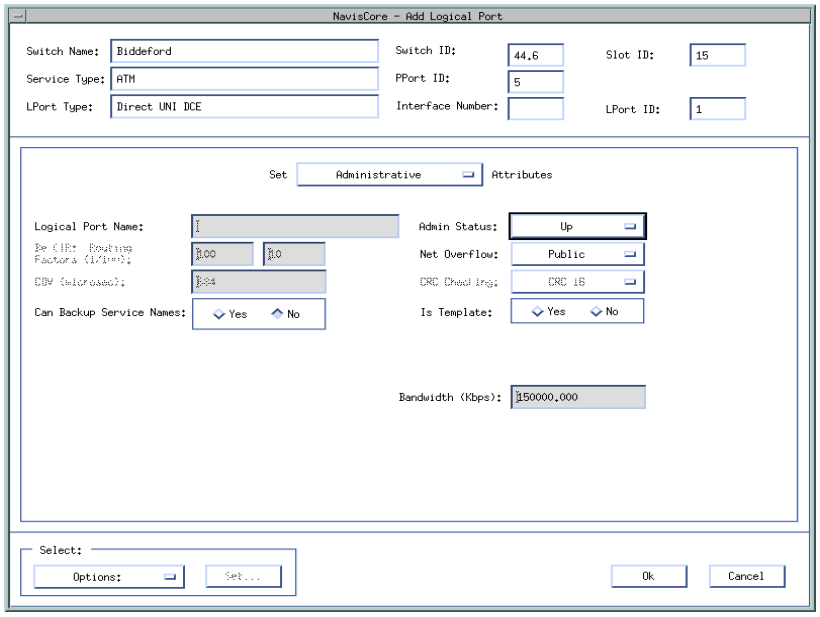

Figure 3-17. Add Logical Port Administrative Attributes Dialog Box

8. Configure the logical port as an ATM UNI DCE port using the configuration instructions in the *NavisCore ATM Configuration Guide.*

9. The next step is to specify the IP interface address for the IP logical port. See "Setting the IP Interface Address" on page 3-11.

IP Server PVCs on the CBX 500

For each IP-enabled ATM logical port, you need a working PVC from an ATM logical port to an IP server logical port.

Creating an IP Server PVC

To create an IP server PVC from the NavisCore menu:

1. Select the appropriate CBX 500 switch icon from the network map.
2. Select Ascend IP Parameters ⇒ Set IP Servers ⇒ Set IP Server PVCs from the Administer menu. The Set All IP Server PVCs on Map dialog box appears (Figure 3-18).

   The Set All IP Server PVCs on Map dialog box initially displays no defined circuit names. To display all of the defined circuit names, position the cursor in the *Search by Name* field and press Enter. This search may take several minutes depending on your configuration.

   For a partial search, enter the selected search criteria in the *Search by Name* field. To use a wildcard search to find a specific circuit name, you can:

   - Use an * to match any number of characters
   - Use a ? to match a single character
   - Use a * to match the * character
   - Use a \? to match the ? character
   - Use a \\ to match the \ character

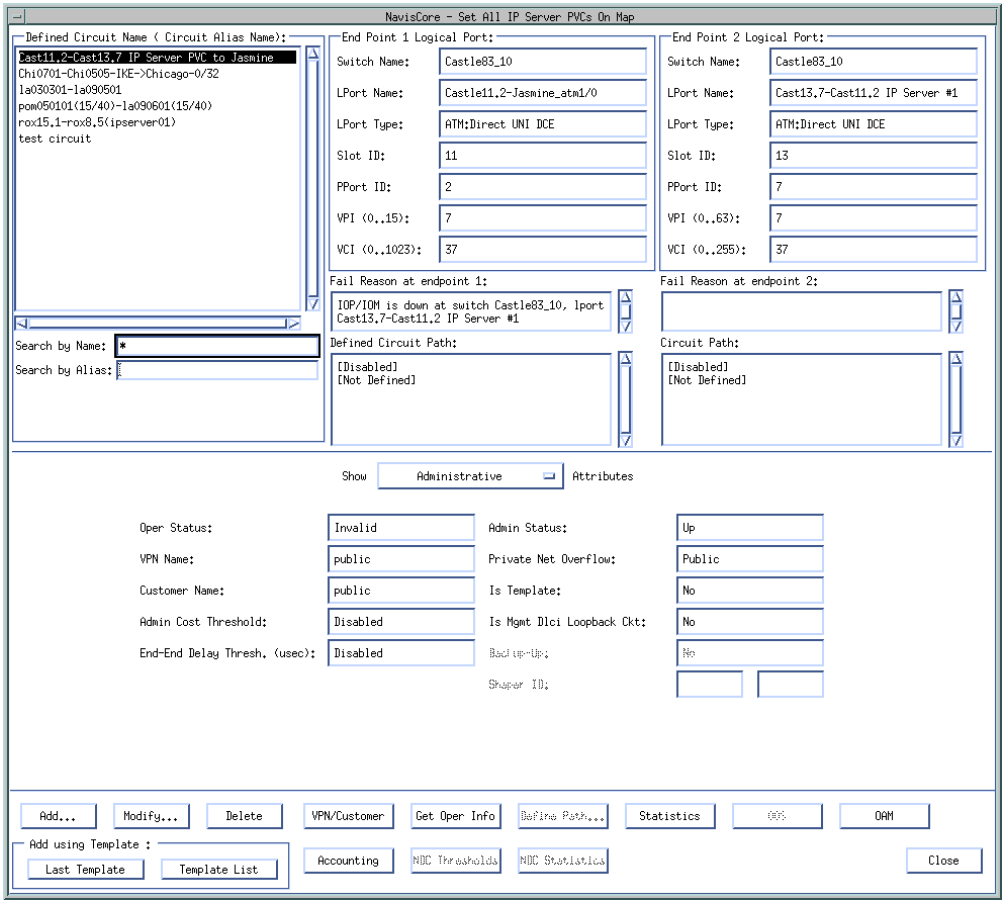

Figure 3-18. Set All IP Server PVCs on Map Dialog Box

3. Choose Add. The Select End Logical Ports dialog box appears.

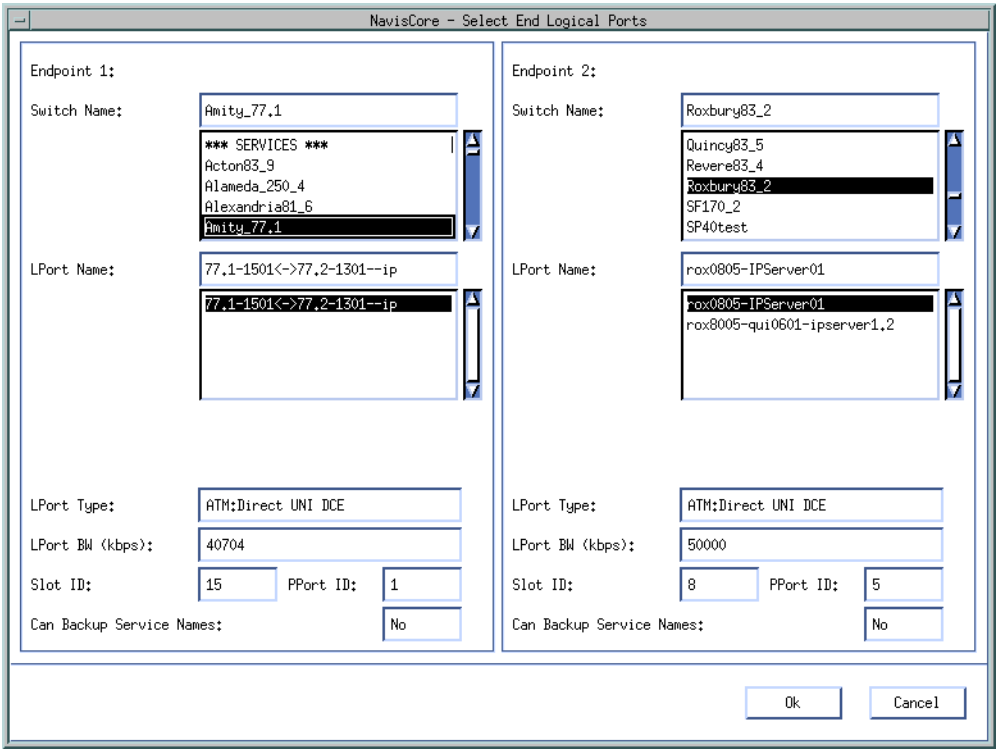

Figure 3-19. Select End Logical Ports

4. Select the name of the switch where Endpoint 1 resides, then select the name of the switch where Endpoint 2 resides.

- Endpoint 1 is the ATM cell endpoint.
- Endpoint 2 is the IP Server endpoint.

In order to accept or transmit IP traffic on a cell-based card you must configure a PVC connection from the cell-based card to an IP server card for all traffic entering the IP port.

5. Select the name of the logical port for Endpoint 1, then select the name of the logical port for Endpoint 2.The Select End Logical Ports dialog box displays information for both Endpoint 1 and Endpoint 2. Table 3-8 describes each of these displayed fields.

Table 3-8. Information Displayed for Endpoints 1 and 2

Field	Action/Description
LPort Type	Displays the logical port type for the selected logical ports.
LPort Bandwidth	Displays the logical port bandwidth for the selected logical ports. At each endpoint, logical ports may have a different bandwidth.
Slot ID	Displays the I/O slot (number) where the IOMs for the selected logical ports reside.
PPort ID	Displays the port ID numbers for the selected logical ports.
Can Backup Service Names	Displays either Yes or No to specify whether or not this logical port can be backed up to a service name binding.

6. Choose OK from the Select End Logical Ports dialog box. The Add PVC dialog box appears.
7. See the *NavisCore ATM Configuration Guide* to define the following attributes for each PVC:

   Administrative

   Traffic Type

   User Preference

   Frame Discard

4

Configuring IP Packet Filters

This chapter describes the following tasks:

- Configuring IP packet filters
- Assigning IP packet filters to a logical port
- Assigning IP packet filters to the host
- Assigning IP packet filters to a circuit
- Viewing an IP packet filter configuration

About Packet Filters

Packet filtering enables a switch to accept or reject inbound or outbound packets by comparing a packet's IP upper-layer header information (see below for the IP header fields) to configured parameters called *filters*, which you define in NavisCore.

You define packet filters based on the following fields in the IP packet header:

- IP Header
- UDP/TCP Header

The following sections describe each of these fields.

IP Header

Source Address — The source address field contains the IP address that sends the packet.

Destination Address — The destination address field contains the IP address that receives the packet.

Type of Service (TOS) — The TOS field indicates the packet's priority.

Protocol — The transport field specifies the protocol (TCP or UDP) that enables the packet to be delivered to the correct destination protocol.

UDP/TCP Header

Source Port — This field contains the 16-bit protocol port number used to demultiplex datagrams among processes waiting to receive them. The source port is optional. When used, it specifies the port to which replies should be sent. If not used, the field should be left blank.

Destination Port — This field contains the 16-bit protocol port number used to demultiplex datagrams among processes waiting to receive them.

For inbound filters, when a packet is received, the forwarding code checks the packet against the interface's list of filters. If the packet matches a filter in the filter list, the packet is accepted or rejected and further filtering is terminated. The packet goes through a similar process for outbound filters, however, the process occurs only after the packet is received and routed to an interface.

Configuring IP Packet Filters

When you define an IP packet filter, you specify specific parameters that control the processing of inbound and/or outbound packets. After you define the filter, you can assign it to IP logical ports, the switch itself (host), or PVCs.

This section describes how to:

- Define an IP packet filter
- Assign an IP packet filter to a logical port
- Assign an IP packet filter to a host (switch)
- Assign an IP packet filter to a circuit
- View an IP packet filter's configuration and its associated logical port and/or circuit

You can create a maximum of 1024 packet filters per switch.

You can define 128 logical port/circuit filter bindings per IOP.

You can assign a maximum of 32 inbound and 32 outbound filters per logical port.

You can assign a maximum of 32 inbound and 32 outbound filters per circuit.

Defining an IP Packet Filter

To define an IP packet filter:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (Figure 4-1).

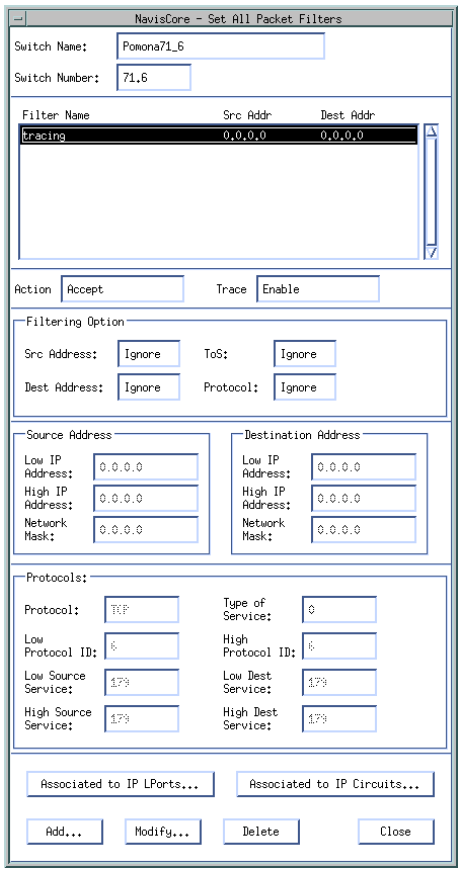

Figure 4-1. Set All Packet Filters Dialog Box

The Set All Packet Filters dialog box displays the following buttons.

Button	Function
Associated to IP LPorts	Displays the logical ports that are using a selected packet filter. For more information, see "Viewing an IP Packet Filter's Configuration" on page 4-22.
Associated to IP Circuits	Displays the circuits that are using a selected packet filter. For more information, see "Viewing an IP Packet Filter's Configuration" on page 4-22.
Add	Enables you to add a filter.
Modify	Enables you to modify an existing filter.
Delete	Enables you to delete an existing filter.

3. Choose Add. The Set Filter dialog box appears (Figure 4-2).

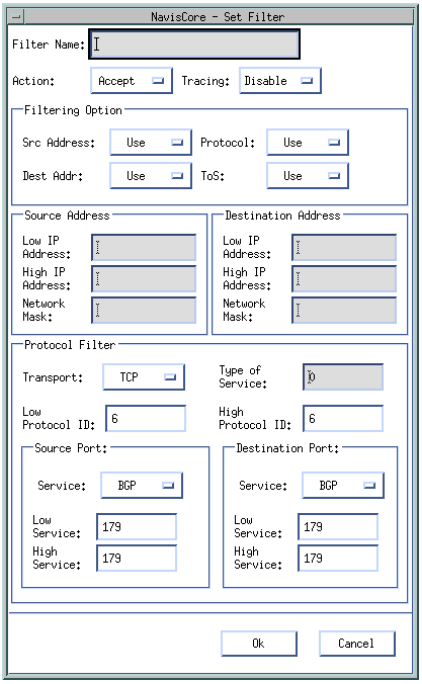

Figure 4-2. Set Filter Dialog Box

4. Complete the fields as described in Table 4-1.

Table 4-1. Set Filter Fields

Field	Action/Description
Filter Name	Enter a filter name to identify the filter.
Action	Select one of the following options: *Accept* – This parameter instructs the switch to accept packets that match the filtering criteria. *Reject* – This parameter instructs the switch to reject packets that match the filtering criteria.
Tracing	Select one of the following options: *Enable* – This parameter instructs the switch to pass matched packets to the trace manager. *Disable* – This parameter instructs the switch not to pass matched packets to the trace manager.
Filtering Option	
Src Address	Select one of the following options: *Use* – To filter packets based on the source address field in the IP packet header. *Ignore* – To ignore filtering based on the source address field in the IP packet header. If you choose Ignore, the source address fields are grayed out and cannot be defined.
Protocol	Select one of the following options: *Use* – To filter packets based on the source address field in the IP packet header. *Ignore* – To ignore filtering based on the source address field in the IP packet header. If you choose Ignore, the protocol fields are grayed out and cannot be defined.
Dest Addr	Select one of the following options: *Use* – To filter packets based on the destination address field in the IP packet header. *Ignore* – To ignore filtering based on the destination address field in the IP packet header. If you choose Ignore, the Destination Address fields are grayed out and cannot be defined.
ToS	Select one of the following options: *Use* – To filter packets based on the destination address field in the IP packet header. *Ignore* – To ignore filtering based on the Type of Service field in the IP packet header. If you choose Ignore, the Type of Service field is grayed out and cannot be defined.

Table 4-1. Set Filter Fields (Continued)

Field	Action/Description
Source Address	
Low IP Address	Enter the low IP address of the node that sends the packet. When you specify the source address you specify one IP address (in the Low IP Address field) or you can specify a range between the lowest and highest IP address. If a packet's source address is within the range, there is a match. ***Note:*** *If you want to filter packets coming from one IP address, specify the IP address in the low IP address field. You do not have to specify a value in the high IP address field.*
High IP Address	Enter the high IP address of the node that sends the packet (the default is *high IP address=low IP address).*
Network Mask	Enter the Network Mask that applies to the source address.
Destination Address	
Low IP Address	Enter the low IP address of the node that receives the packet. When you specify the destination address you specify one IP address (in the Low IP Address field) or you can specify a range between the lowest and highest IP address. If a packet's source address is within the range, there is a match.
High IP Address	Enter the high IP address of the node that receives the packet (the default is *high IP address=low IP address).*
Network Mask	Enter the Network Mask that applies to the destination address.

Table 4-1. Set Filter Fields (Continued)

Field	Action/Description
Protocol Filter	
Transport	Select the packet's transport protocol type: *TCP* – Transmission Control Protocol. *UDP* – User Datagram Protocol *Others* – You must specify protocol IDs in the low and high protocol ID fields. Transport refers to the protocol (TCP, UDP, or Others) that enables the packet to be delivered to the correct destination protocol. ***Note:*** *When you select TCP or UDP, the low and high protocol fields are automatically filled in with the protocol's corresponding protocol ID.* *In addition, if you select TCP or UDP, you must specify the source and destination port fields. However, if you select Others, the source and destination port sections are grayed out and cannot be defined.*
Type of Service	Enter a value between 0 and 254. Protocols use the type of service value to specify the packet's priority.
Low Protocol ID	If you selected Others in the Transport field, enter the low protocol ID. See RFC 1700 for protocol ID numbers. You can either specify *one value* (in this field) or you can enter a range between this value and the high protocol ID. If the packet's protocol ID is between the low and high protocol ID, there is a match.
High Protocol ID	If you selected Others in the Transport field, enter the high protocol ID. See RFC 1700 for protocol ID numbers. When you enter this value, you enter a range between the low protocol ID and this value. If the packet's protocol ID is between the low and high protocol ID, there is a match.

Table 4-1. Set Filter Fields (Continued)

Field	Action/Description
Source Port	
Service	*If you selected TCP* in the transport protocol field, select one of the following protocols: *BGP* – Border Gateway Protocol. *FTP* – File Transfer Protocol. *Gopher* – Protocol that facilitates internet access. *IRC* – Internet Relay Chat Protocol. *Talk* – Unit talk application. *Telnet* – Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection. *WWW* – World Wide Web. *Ignore* – Enables you to filter on all UDP packets. *Other* – You must specify port numbers in the low and high service fields. *If you selected UDP* in the Transport Protocol field, select one of the following protocols: *RIP* – Routing Information Protocol. *SNMP* – Simple Network Management Protocol. *Traps (SNMP)* – Message sent by an SNMP agent to an NMS station to indicate that an event occurred. *TFTP* – Trivial File Transfer Protocol. *Ignore* – Enables you to filter on any service. *Other* – You must specify port numbers in the low and high service fields. ***Note:*** *When you select a service, the low and high service fields in the source port field are automatically filled in with the service's corresponding port number.*
Low Service	If you selected Other in the Service field, enter the low service port number. See RFC 1700 for the port numbers. ***Note:*** *To filter packets that have the same service port number, specify the port number in the low service field. You do not have to specify a value in the high service field.*

Table 4-1. Set Filter Fields (Continued)

Field	Action/Description
High Service	If you selected Other in the Service field, enter the high service port number. See RFC 1700 for the port numbers. When you enter this value, you enter a range between the low service port number and this value. If the packet's service port number is between the low and high service port numbers, there is a match.
Destination Port	
Service	See "Service" field description in the source port field section.
Low Service	See "Low Service" field description in the source port field section.
High Service	See "High Service" field description in the source port field section.

5. Choose OK.
6. At the Set IP Filter List dialog box, choose Close.

Packet Filter Configuration Example

The following configuration is an example of a filter that restricts packets coming from a specified source IP address.

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (Figure 4-1 on page 4-4).
3. Choose Add. The Set Filter dialog box appears (Figure 4-2 on page 4-5).
4. In the Filter Name field, enter:

   `reject152.148.51.118`
5. In the Action field, select Reject.
6. In the Tracing field, select Disable to disable the trace manager.
7. In the Filtering Option fields;
   - Select Use in the Source Address field.
   - Select Ignore in the Destination Address, Protocol, and TOS fields.

When you select Ignore in these fields, the Protocol Filter, Source Port, and Destination Port fields are grayed out. These fields are disabled and are not used to filter packets.

8. In the Low IP Address field for the source address section, enter:

   `152.148.51.118`

You do not have to specify the high IP address for the source address because you are restricting packets coming from one IP address. However, if you want to restrict packets coming from a range of IP addresses, specify both the low and high IP addresses.

9. In the Network Mask field, enter:

   `255.255.255.255`

   Figure 4-3 displays the specified fields.

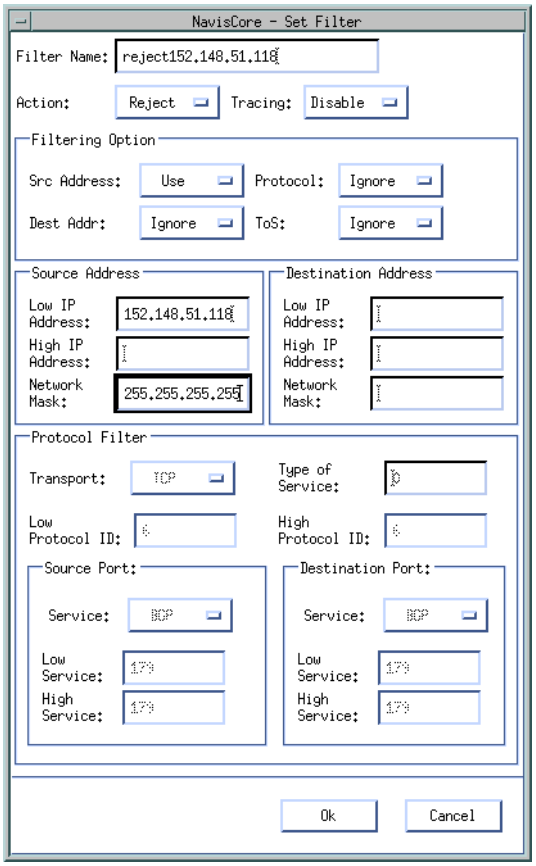

Specify 152.148.51.118 in the Low IP Address field and 255.255.255.255 in the Network Mask field.

The Protocol Filter, TOS, and Source and Destination Port fields are grayed out because you selected Ignore in the Filtering Option fields

Figure 4-3. Example Packet Filter Settings

10. Choose OK.

When you assign this filter to a specific logical port or host, all packets coming from *152.148.51.118* are not allowed to pass through.

Assigning IP Packet Filters to Logical Ports

To assign an IP packet filter to a logical port:

1. Select the switch icon from the network map.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Packet Filters ⇒ Set All Logical Port Filters. The Set All Logical Port Filters dialog box appears (Figure 4-4).

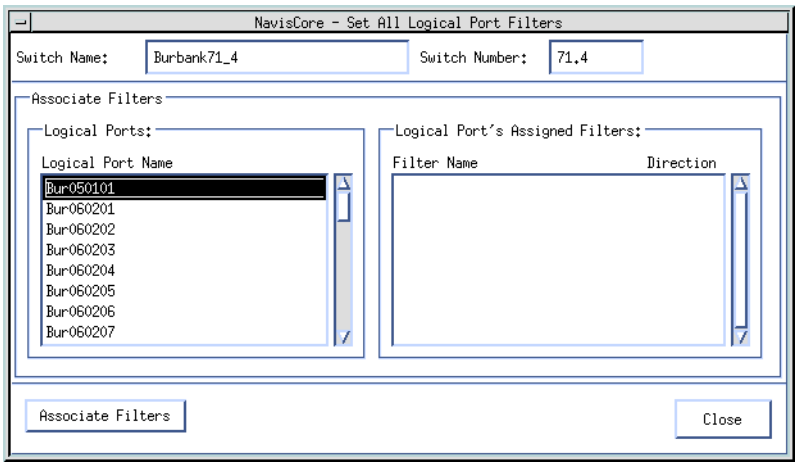

Figure 4-4. Set All Logical Port Filters

3. In the Logical Ports list box, select the logical port with which you want to associate a filter.
4. Choose the Associate Filters button. The Assign Logical Port IP Filter dialog box appears (Figure 4-5).

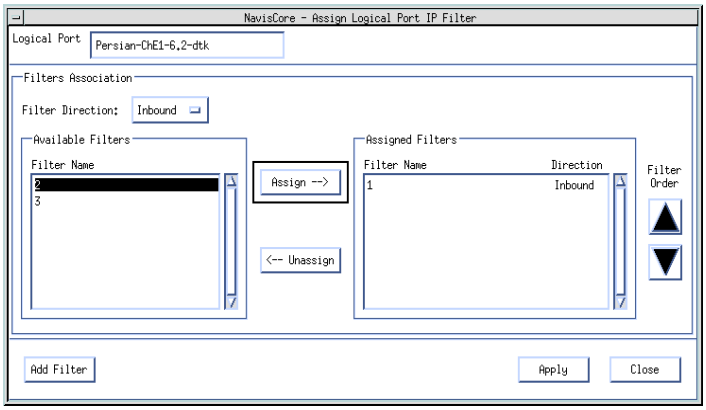

You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filter fields.

Figure 4-5. Assign Logical Port IP Filter Dialog Box

You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Show IP Filter Configuration dialog box appears with the filter's configuration.

Table 4-2 describes the fields on the Assign Logical Port IP Filter dialog box.

Table 4-2. Assign Logical Port IP Filter Fields

Field	Description
Logical Port	Displays the name of the logical port.
Filter Direction	Enables you to indicate the direction (inbound or outbound) in which you want the packets filtered through this logical port.
Available Filters	
Filter Name	The name that identifies the filter.
Assigned Filters	
Filter Name	The name that identifies the filter.
Direction	The direction (inbound or outbound) in which you want the packets filtered.

The following list describes the Assign Logical Port IP Filter buttons.

Button	Function
Assign	Enables you to assign an IP packet filter to this logical port.
Unassign	Enables you to remove an IP packet filter from this logical port.
Filter Order	Enables you to specify the order in which the defined packet filters are applied. When a match occurs, the filtering process ends.
Add Filter	Enables you to configure an additional IP packet filter.
Apply	Applies any of the changes that you have made on this dialog box.

5. In the Filter Direction field, select either Inbound or Outbound to indicate the direction you want the packets filtered through this logical port.
6. From the Available Filters list box, select the filter and choose Assign to assign the IP packet filter to this logical port.
7. Repeat step 6 until you have assigned all the necessary IP packet filters to this logical port.
8. When you are done, choose apply.
9. To configure an additional IP packet filter, choose Add Filter. See "Defining an IP Packet Filter" on page 4-3 for more information.

Assigning IP Filters to the Host (Switch)

You can assign a maximum of 32 IP packet filters per host.

To assign a filter to the host (switch):

1. Select the switch icon from the network map.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Packet Filters ⇒ Set All Host Filters. The Set All Host filters dialog box appears (Figure 4-6).

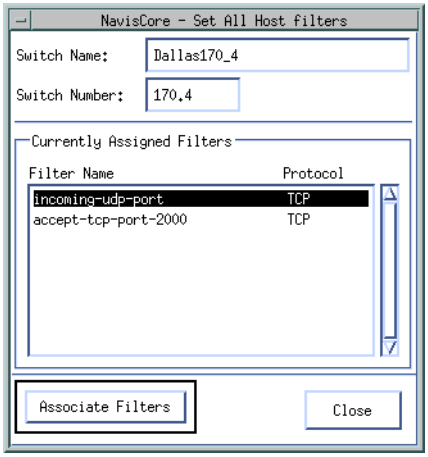

Figure 4-6. Set All Host filters Dialog Box

3. Choose Associate Filters. The Associate Host Filters dialog box appears (Figure 4-7).

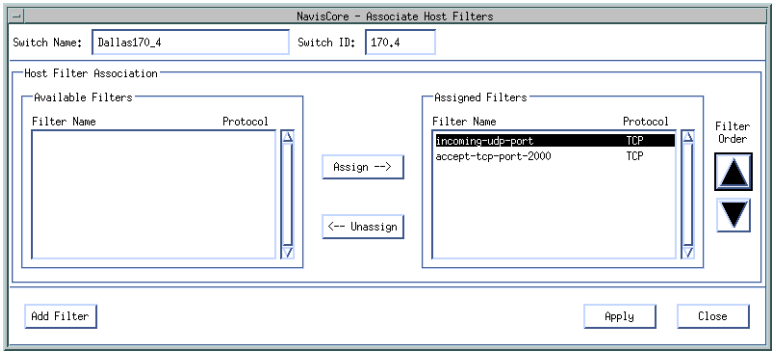

The Filter Order buttons enable you to change the order in which filters are applied.

Figure 4-7. Associate Host Filters Dialog Box

You can view a packet filter's configuration by double-clicking on the desired filter in either the Available or Assigned Filters fields. The Show IP Filter Configuration dialog box appears with the filter's configuration.

Table 4-3 describes the fields on the Associate Host Filters dialog box.

Table 4-3. Associate Host Filters Fields

Field	Description
Switch Name	Displays the name of the switch.
Switch ID	Displays the switch ID.
Available Filters	
Filter Name	The name that identifies each of the defined filters.
Protocol	Displays TCP, UDP, or Others to indicate the filter's transport protocol. See Table 4-1 for a description of each of these protocol types.
Assigned Filters	
Filter Name	The name that identifies each of the defined filters that are assigned to the switch.
Protocol	Displays TCP, UDP, or Others to indicate the filter's transport protocol. See Table 4-1 for a description of each of these protocol types.

The Associate Host Filters dialog box provides the following buttons:

Button	Function
Assign	Enables you to assign an IP packet filter to this host.
Unassign	Enables you to delete an IP packet filter from this host.
Filter Order	Enables you to specify the order in which the defined packet filters are applied. When a match occurs, the filtering process terminates.
Add Filter	Enables you to configure an additional IP packet filter.
Apply	Applies any of the changes that you have made on this dialog box.

4. From the Available Filters List box, select the filter and choose Assign to assign the IP packet filter to this switch.

5. Repeat step 4 until you have added the necessary IP packet filters to this switch.
6. When you are done, choose Apply.
7. To configure an additional IP packet filter, choose Add Filter. See "Defining an IP Packet Filter" on page 4-3 for more information.

Assigning IP Packet Filters to Circuits

Circuit filters are similar to logical port filters but differ in that you apply circuit filters to individual DLCIs (for Frame Relay circuits) or individual VPIs/VCIs (for ATM circuits). Before you assign circuit filters to PVCs, you must define these PVCs. For more information, refer to Chapter 3, "Configuring IP Logical Ports and IP Servers".

If you assign packet filters to both logical ports and circuits, the order in which packets are filtered are as follows:
Inbound
Circuit Filters⇒ Logical Port Filters
Outbound
Logical Port Filters⇒ Circuit Filters

To assign a filter to the circuit:

1. Select the switch icon from the network map.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Packet Filters ⇒ Set All Circuit Filters. The Set All IP Circuit Filters dialog box appears.

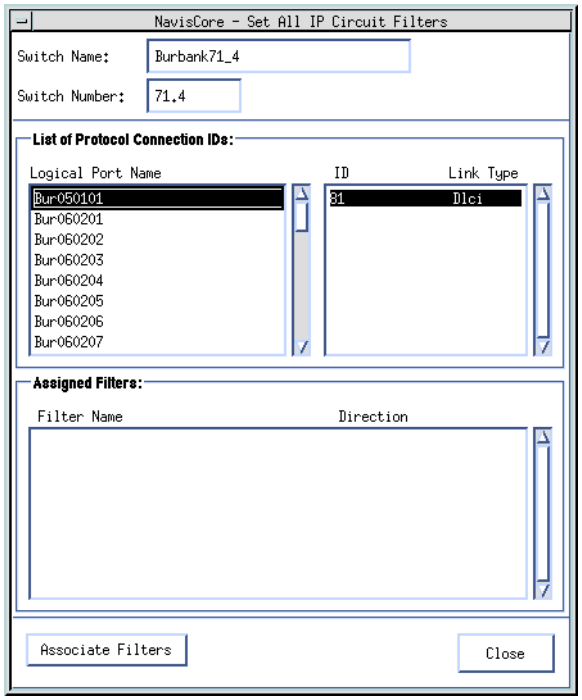

Figure 4-8. Set All IP Circuit Filters Dialog Box

3. Choose Associate Filters. The Associate IP Circuit Filter List dialog box appears.

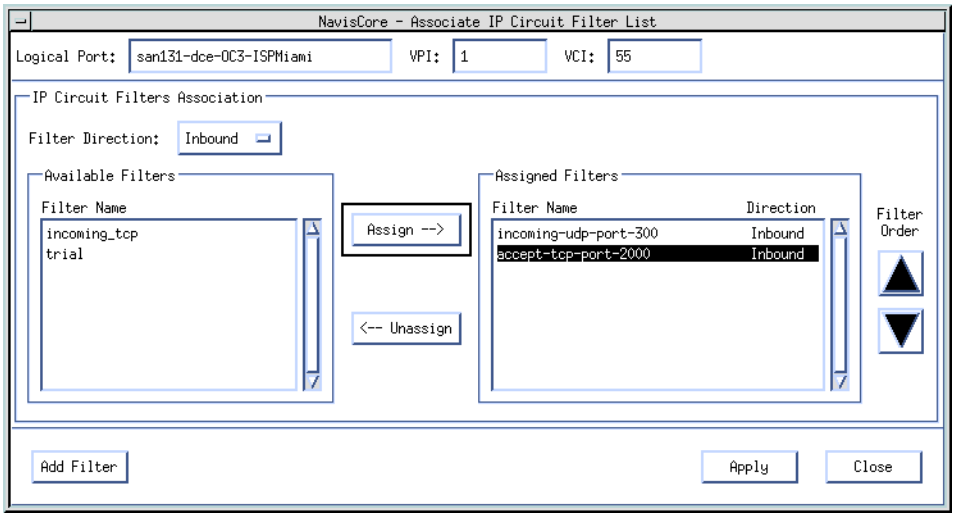

Figure 4-9. Associate IP Circuit Filter List

You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Show IP Filter Configuration dialog box appears with the filter's configuration.

Table 4-4 describes the Associate IP Circuit Filter List dialog box fields.

Table 4-4. Associate IP Circuit Filter List Fields

Field	Description
Logical Port	Displays the circuit name.
VPI/VCI (For ATM logical ports)	Displays the circuit's VPI/VCI.
DLCI (For Frame Relay logical ports)	Displays the circuit's DLCI.
Filter Direction	Enables you to indicate the direction (inbound or outbound) in which you want the packets filtered through this circuit.
Available Filters	
Filter Name	The name that identifies the filters available to this circuit.
Assigned Filters	
Filter Name	The name that identifies the filter(s) assigned to this circuit.
Direction	Indicates the direction (inbound or outbound) in which you want the packets filtered.

The Associate Protocol Filter dialog box provides the following buttons:

Button	Function
Assign	Enables you to assign an IP packet filter to this host.
Unassign	Enables you to delete an IP packet filter from this circuit.
Filter Order	Enables you to specify the order in which the defined packet filters are applied. When a match occurs, the filtering process terminates.
Add Filter	Enables you to configure an additional IP packet filter.
Apply	Applies any of the changes that you have made on this dialog box.

4. From the Available Filters List box, select the filter and choose Assign to assign the IP packet filter to this switch.
5. Repeat step 4 until you have added the necessary IP packet filters to this switch.
6. When you are done, choose Apply.
7. To configure an additional IP packet filter, choose Add Filter. See "Defining an IP Packet Filter" on page 4-3 for more information.

Viewing an IP Packet Filter's Configuration

Once you define an IP packet filter and associate it, you can view its configuration and associated logical port or circuit. Use the following steps to view an IP packet filter configuration for a logical port:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (Figure 4-10).

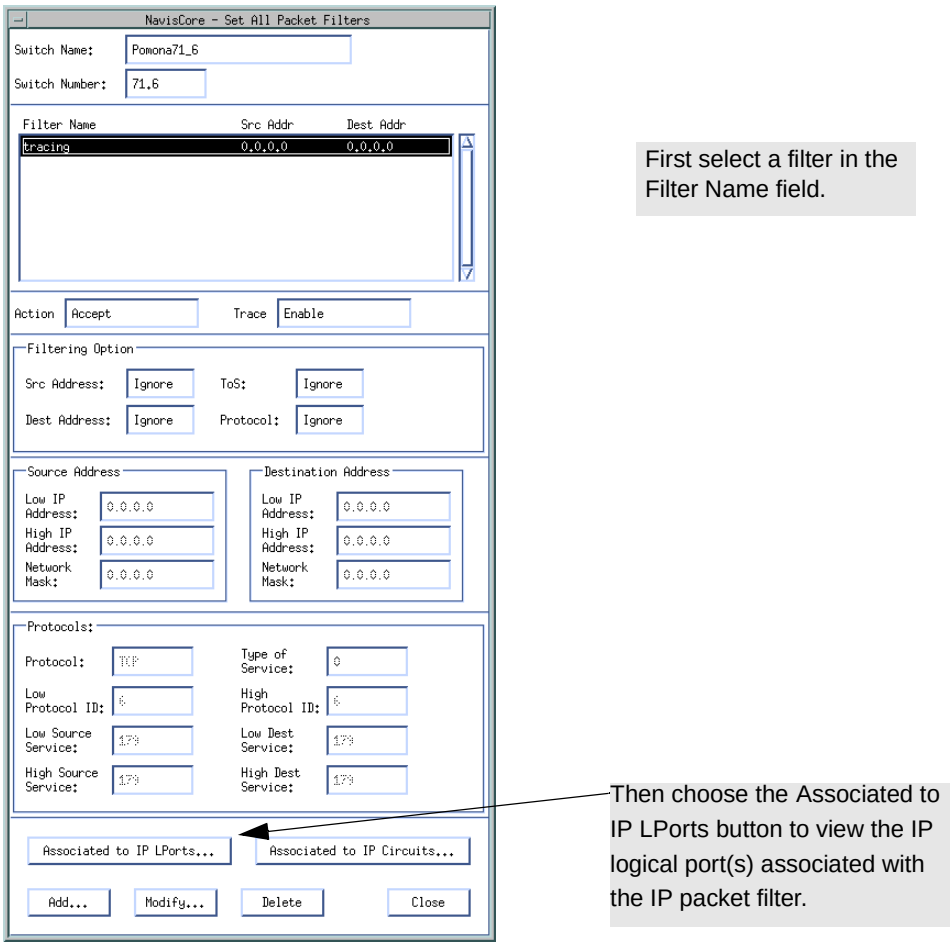

Figure 4-10. Set All Packet Filters Dialog Box

3. Do the following:

   a. Select a filter in the filter field.

   b. Choose the Associated to IP LPorts button. The Logical ports using the Packet Filter dialog box appears (Figure 4-11).

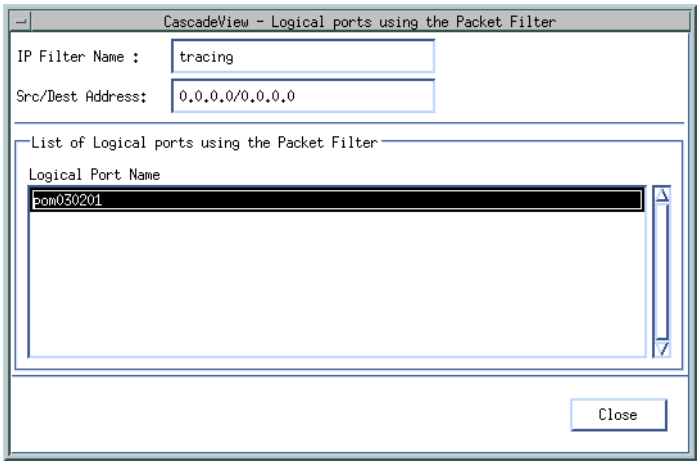

Figure 4-11. Logical ports using the Packet Filter Dialog Box

If you selected a packet filter that is not assigned to an IP logical port, the following message appears:

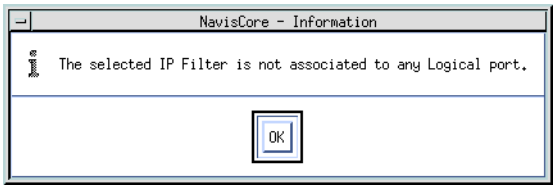

Figure 4-12. Packet Filter Error Message Dialog Box

4. Choose Close.

5. To view a packet filter that is assigned to a circuit, perform step 1 through step 3, except choose the **Associated to IP circuits** button instead of the **Associated to IP LPorts** button.

5

Provisioning IP Quality of Service

This chapter describes the following tasks:

- Configuring an IP QoS PVC
- Defining an IP Flow Profile
- Assigning an IP Flow Profile to a Logical Port

About IP Flow Profiles

An IP Flow Profile directs traffic from a source address to a destination address over an IP QoS PVC. When a logical port receives an IP datagram, it compares it with a predefined IP Flow Profile. If there is a match, the logical port forwards the datagram over the specified IP QoS PVC. If the datagram does not match an IP Flow Profile, the logical port forwards it over a MultiPoint-to-Point Tunnel (MPT), which is a best effort PVC. You associate IP Flow Profiles with existing PVCs which you set up with Quality of Service attributes (QoS class, traffic descriptors such as PCR, SCR).

Some important facts about IP QoS and flow profiles:

- One PVC may be shared by multiple flow profiles.
- One flow profile may be shared by multiple logical ports.
- You can assign up to 100 flow profiles to one switch.
- You can assign up to 50 flow profiles to one logical port.
- You can modify an IP address in a flow profile with a CIDR mask. For example, if the address field is set to 1.2.3.4 with a mask of 255.255.255.0, an address of 1.2.3. + anything causes a match.
- You must verify there is enough bandwidth on the PVC for the IP profiles that share the PVC.

- You configure IP filtering before IP QoS.
- You can order the importance of the IP Flow Profiles. The first match wins.
- When you reboot the switch, the PVC takes time to become active.
- The CP maintains the QoS information, but then distributes it to the appropriate IOPs so the CP is not required for proper QoS operation. This step eliminates a single point of failure.

About IP QoS PVCs

The IP QoS PVC is like an ordinary PVC, except it is a switch-to-switch PVC, not logical port-to-logical port. You can configure an infinite number of PVCs between switches. The IP QoS PVC is defined by its logical port ID and source DLCI. To create the IP QoS PVC, you use a point-to-point connection. The endpoints of the QoS PVC are special logical ports that you configure when you add the first node to the map. This release supports IP QoS Frame Relay circuits only.

Provisioning QoS

This section describes how to:

- Configure an IP QoS PVC
- Define an IP flow profile and attach it to the IP QoS PVC
- Assign an IP flow profile to a logical port

Configuring an IP QoS PVC

To configure an IP QoS PVC, you first create the circuit endpoints, then define and name the circuit connection. You also assign administrative and user preference attributes to the circuit.

To configure an IP QoS PVC:

1. From the Administer menu, select Ascend IP Parameters ⇒ Set IP QoS PVCs. The Set All IP QoS PVCs On Map dialog box appears (Figure 5-1).

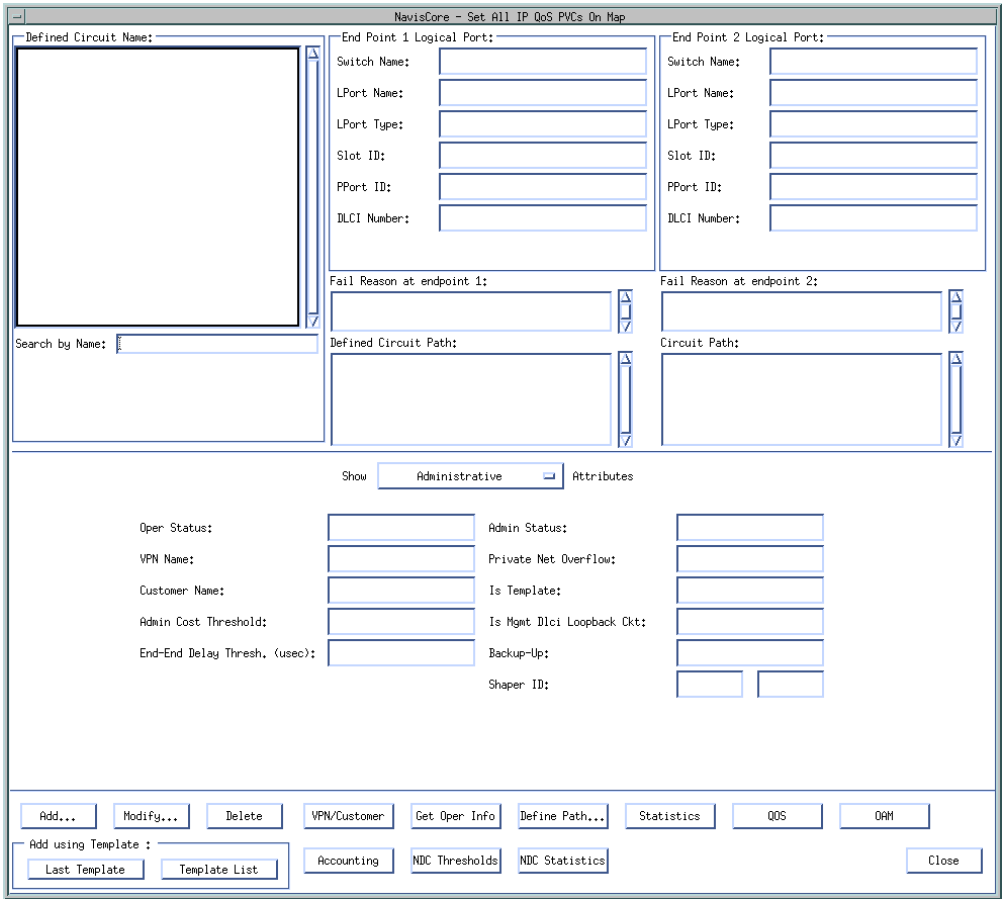

Figure 5-1. Set All IP QoS PVCs On Map Dialog Box

The Set All IP QoS PVCs on Map dialog box initially displays no defined circuit names. To display all of the defined circuit names, position the cursor in the *Search by Name* field and press Enter. This search may take several minutes depending on your configuration.

For a partial search, enter the selected search criteria in the *Search by Name* field. To use a wildcard search to find a specific circuit name, you can:

- Use an * to match any number of characters
- Use a ? to match a single character
- Use a * to match the * character
- Use a \? to match the ? character
- Use a \\ to match the \ character

Table 5-1 describes the buttons on the Set All PVCs dialog box.

Table 5-1. Set All IP QoS PVCs on Map Buttons

Command Button	Description
Add/Modify/Delete	Enables you to add a new circuit or Modify or Delete an existing circuit. ***Note**: If the PVC loopback status field does not display NONE, do not attempt to modify or delete the selected circuit.*
VPN/Customer	Displays the virtual private network customer's name.
Get Oper Info	Displays a status message in the *Oper Status* field about the selected circuit. For more information, see the *NavisCore Diagnostic and Troubleshooting Guide*.
Define Path	Enables you to manually define a circuit path.
Statistics	Displays the summary statistics for the selected circuit. For information, see the *NavisCore Diagnostic and Troubleshooting Guide*.
QoS	Displays the Quality of Service values for the selected circuit.
OAM Alarms (*ATM CS and IWU modules only*)	Displays the OAM alarms which indicate whether the circuit is up or down. These alarms send a signal to the logical port whenever the circuit goes down or comes back up.
Add Using Template	If you have already defined a circuit configuration and saved it as a template, use this option to define a new circuit. Choose *Last Template* to use the last template you defined for this switch. Choose *Template List* to display a list of templates previously defined for this map.
Accounting	Accesses the accounting functions for a PVC. For more information, see the *NavisXtend Accounting Server Administrator's Guide.*
NDC Thresholds	Displays the configured network data collection (NDC) thresholds for the selected circuit. For more information about these thresholds, see the *NavisCore Diagnostic and Troubleshooting Guide.*
NDC Statistics	Displays the NDC statistics for the selected circuit. For more information about NDC statistics, see the *NavisCore Diagnostic and Troubleshooting Guide.*
Close	Exits the dialog box and returns you to the network map.

2. Choose Add. The Select End Logical Ports dialog box appears (Figure 5-2).

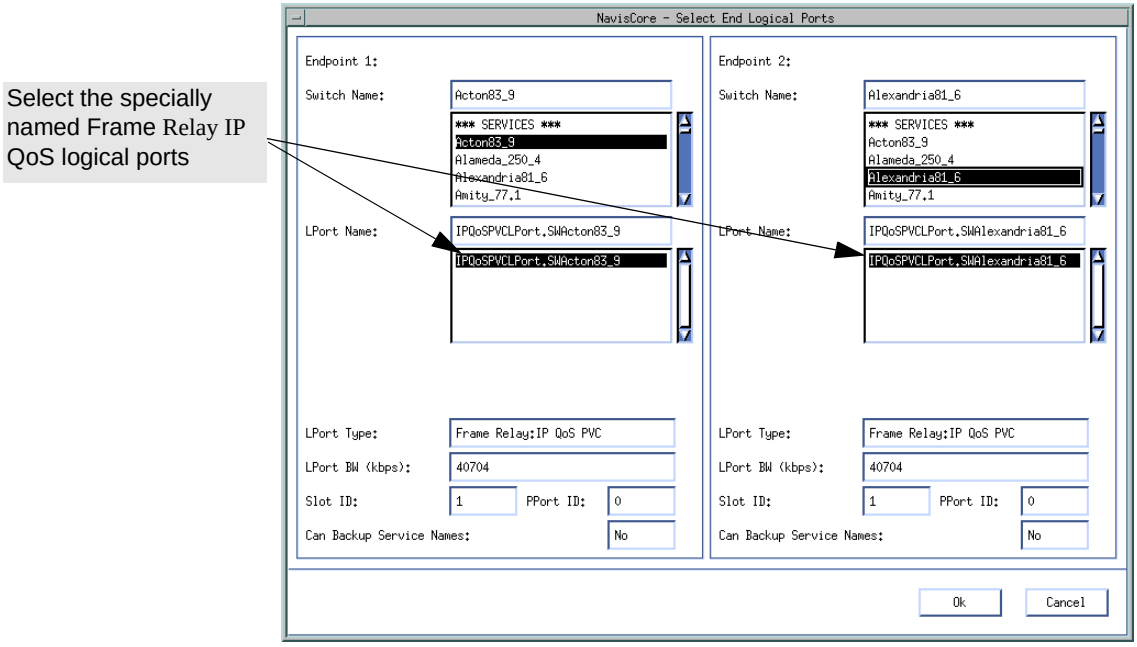

Figure 5-2. Select End Logical Ports Dialog Box

The Select End Logical Ports dialog box displays information based on configuration selections you made. Table 5-2 describes each field.

Table 5-2. Select End Logical Ports Fields

Field	Description
LPort Type	Displays the logical port type for each port in the circuit configuration.
LPort Bandwidth	Displays the bandwidth for each logical port in the trunk configuration.
Slot ID	Displays the I/O slot (number) in which the module resides.
PPort ID	Displays the port number for the physical port.
Can Backup Service Names	Displays either Yes or No to specify whether or not this logical port can be backed up to a service name binding.

3. Configure Endpoint 1 and Endpoint 2 as follows:
   a. Select a switch name from the list.
   b. Select the Frame Relay IP QoS logical port.

4. Choose OK. The Add PVC dialog box appears displaying the current parameters.

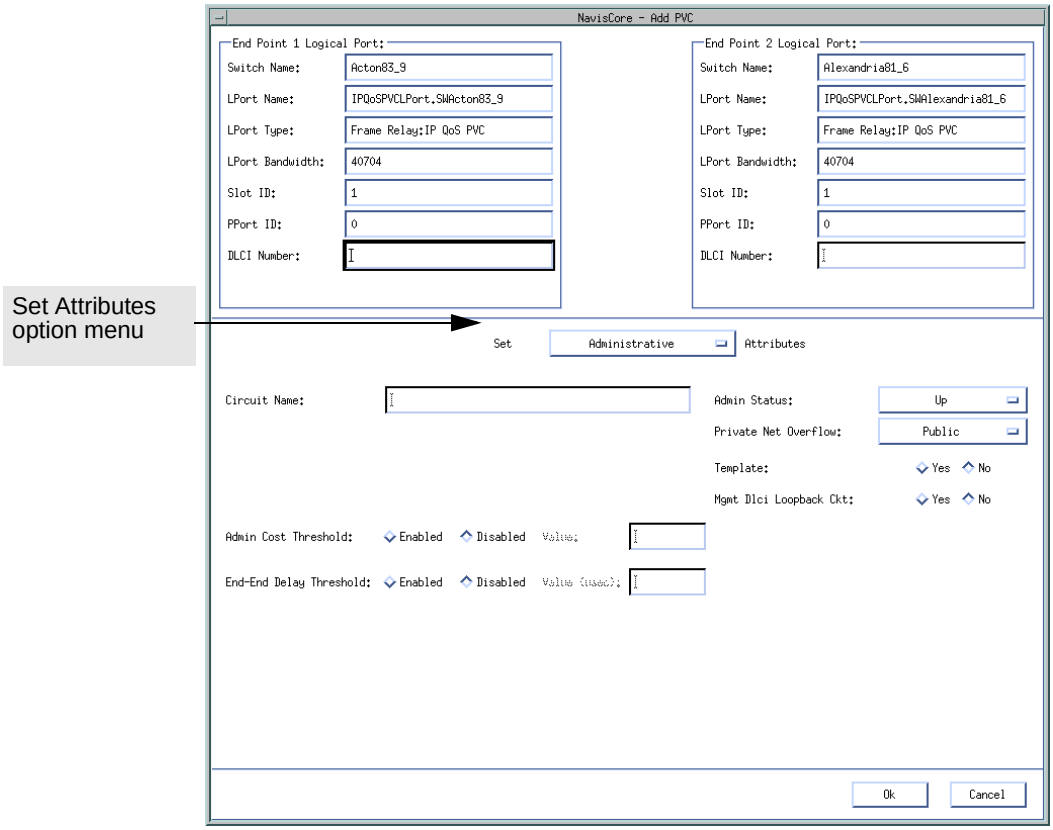

Figure 5-3. Add PVC-Set Administrative Attributes Dialog Box (Frame Relay:IP QoS PVC)

5. Access the Set Attributes option menu and complete the circuit attributes as described in the following sections.

Administrative Attributes

Complete the administrative attributes fields described in Table 5-3.

Table 5-3. Set Administrative Attributes Fields

Field	Action/Description
DLCI Number	Enter a unique DLCI for this logical port. For more information, see "Setting the DLCI for Frame Relay Logical Ports" on page 3-14.
Circuit Name	Enter any unique, continuous, alphanumeric name for the QoS circuit. Do not use parentheses and asterisks.
Admin Status	Select Up or Down to activate or deactivate the circuit. *Up (default)* – Activates the circuit. *Down* – Takes the circuit off-line to run diagnostics such as PVC loopback.
Private Net Overflow	(*For Virtual Private Networks*) Set the Private Net Overflow parameters, which determine whether circuits originating from an Lport will be restricted to trunks of their own VPN or use public (shared) trunks during overflow conditions. Options include: *Public (default)* – Enables the circuit to use public trunks during traffic overflow or trunk failure conditions. *Restrict* – Restricts trunks to their own virtual private network.
Template	(*Optional*) Save these settings as a template to use when configuring another circuit with the same options. To create a template, choose Yes in the *Is Template* field.

Traffic Type Attributes

1. Select Traffic Type from the Set Attributes option menu.

The traffic type fields appear (Figure 5-4).

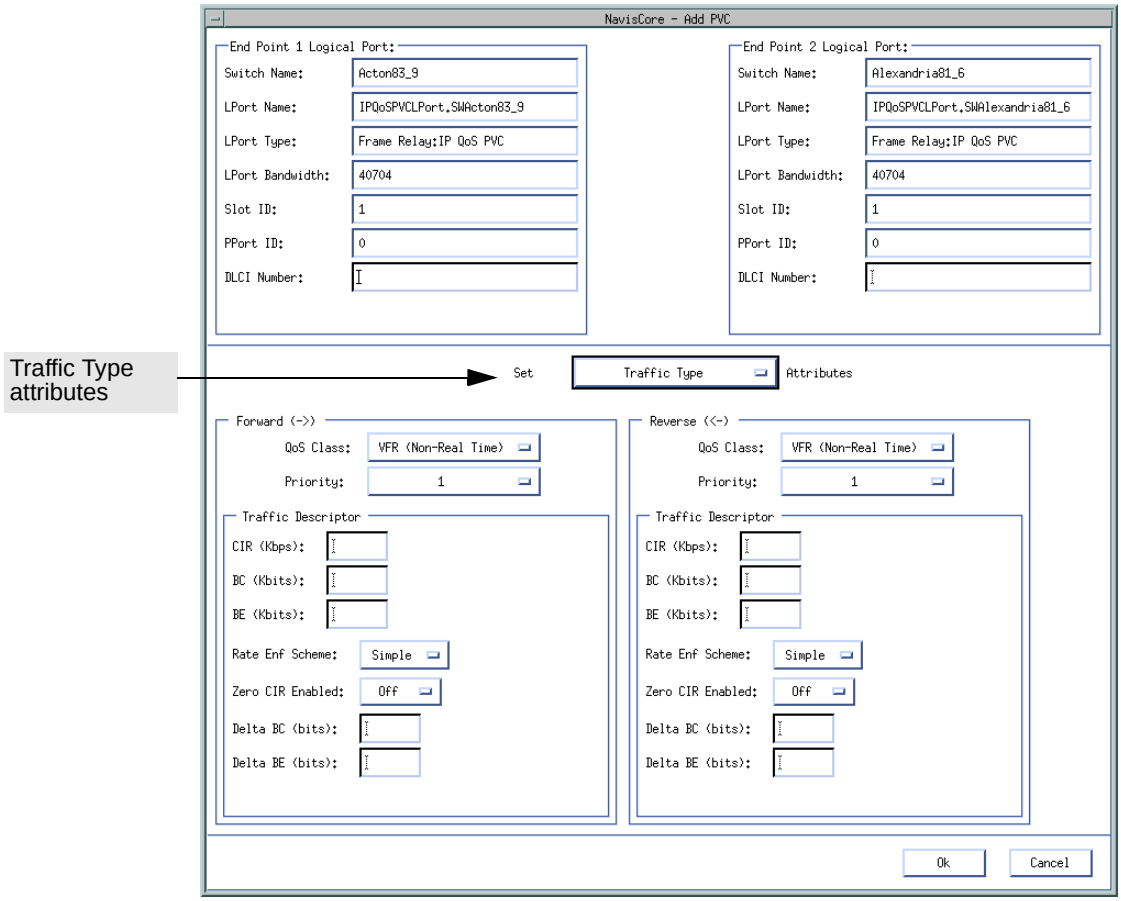

Figure 5-4. Add PVC - Set Traffic Type Attributes Dialog Box (Frame Relay:IP QoS PVC)

2. Complete the required fields described in Table 5-4.

> The left column beneath the (–>) arrow represents the logical port for the circuit that connects Endpoint 1 to Endpoint 2. The right column beneath the (<–) arrow represents the logical port for the circuit that connects Endpoint 2 to Endpoint 1. Enter values in both columns.

Table 5-4. Add PVC - Set Traffic Type Attributes Fields

Field	Action/Description
CIR (Kbps) (Committed Information Rate)	Enter the rate in Kbps at which the network transfers data under normal conditions. Normal conditions refer to a properly designed network with ample bandwidth and switch capacity. The rate is averaged over a minimum increment of the committed rate measurement interval (Tc). The value on each PVC is asymmetric (you can set a different CIR in each direction), which provides more efficient use of bandwidth.
BC (Kbits) (Committed Burst Size)	Enter the maximum amount of data, in Kbits, that the network attempts to transfer data under normal conditions during a specified time interval, Tc. Tc is calculated as BC/CIR. This value must be greater than zero and is typically set to the same value as CIR.
BE (Kbits) (Excess Burst Size)	Enter the maximum amount of uncommitted data, in Kbits, the network will attempt to deliver during a specified time interval, Tc. Tc is calculated as BC/CIR. The network treats this data as “discard eligible” (DE) data.
Circuit Priority (Fwd/Rev)	Select *1*, *2*, or *3* to configure the priority of data being transmitted on this circuit. Circuit priority determines the data’s forwarding priority. The highest priority is 1 (do not discard data); the lowest is 3 (discards data). See the *Networking Services Technology Overview* for information about congestion control and circuit priority.
Zero CIR Enabled (Fwd/Rev)	Set the CIR parameter to *On* or *Off*. *On* – Indicates that the PVC has an assigned CIR value of zero and is a best-effort delivery service. Customer data that is subscribed to zero CIR service can burst to the port speed if there is network bandwidth available to deliver frames. However, no frame delivery guarantees are made. All frames entering the network on zero CIR PVCs have DE set to one. *Off (default)* – Disables zero CIR. ***Note**: If you set Zero CIR Enabled to On you cannot set the CIR, Bc, and Be values.*
Rate Enf Scheme	Select *Simple (default)* or *Jump*. The configurable rate enforcement scheme provides more flexibility, increased rate enforcement accuracy, and improved switch performance. **Note**: *If you select the Simple scheme, the “bad” PVC detection feature is disabled.*

Table 5-4. Add PVC - Set Traffic Type Attributes Fields (Continued)

Field	Action/Description
Delta BC (bits)	Set the number of Delta Bc bits for this circuit between 0 - 65528 (*default* 65528). This value represents the maximum number of bits the network agrees to transfer over the circuit (as committed bits) during the measurement interval, provided there is positive committed bit (Bc) credits before receiving the frame, but negative Bc credits after accepting the frame.
Delta BE (bits)	Set the number of Delta Be bits for this circuit between 0 - 65528 (*default* 65528). This value represents the maximum number of bits the network agrees to transfer over the circuit (as excess bits) during the measurement interval, provided there is positive excess bit (Be) credits before receiving the frame, but negative Be credits after accepting the frame.
Forward QoS Class	Ascend currently supports only the *VBR non-real time* class of service value for non-real time applications. This option enables the circuit to transfer large amounts of data over a long period of time using a pre-established ATM connection. Class of service values enable you to prioritize circuit traffic.
Reverse QoS Class	Ascend currently supports only the *VBR non-real time* class of service value for non-real time applications. This option enables the circuit to transfer large amounts of data over a long period of time using a pre-established ATM connection. Class of service values enable you to prioritize circuit traffic.

User Preference Attributes

1. Select User Preference from the Set Attributes option menu.

The user preference fields appear (Figure 5-5).

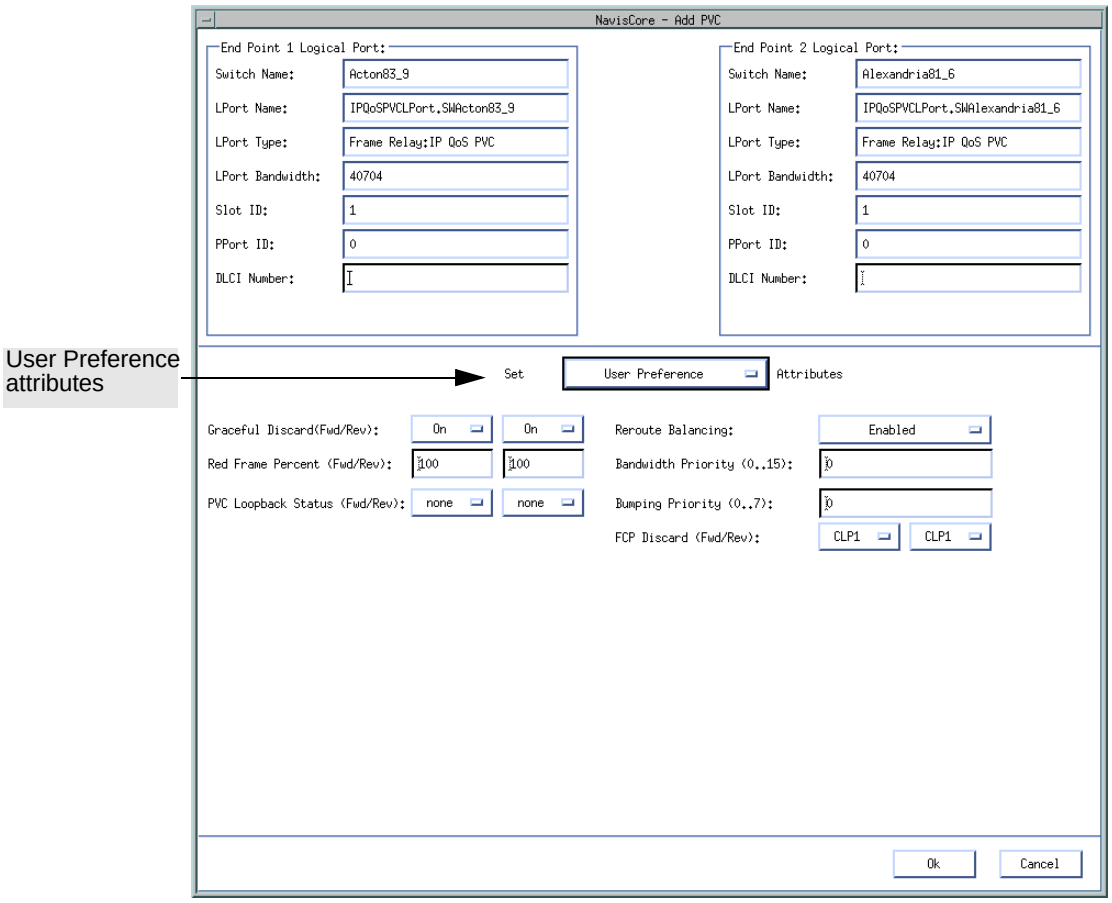

Figure 5-5. Add PVC – Set User Preference Attributes Dialog Box

2. Complete the required fields described in Table 5-5.

Table 5-5. Add PVC - Set User Preference Fields

Field	Action/Description
Graceful Discard (Fwd/Rev)	Select either *On* or *Off* to define how this circuit handles "red" packets. Red packets are designated as those bits received during the current time interval that exceed the committed burst size (Bc) and excess burst size (Be) thresholds, including the current frame. The Discard Eligible (De) bit for a red packet is set to 1, meaning the network can discard this packet unless Graceful Discard is set to On. *On* – Forwards some red packets if there is no congestion. *Off* – Immediately discards red packets.
Red Frame Percent (Fwd/Rev)	Set this value only if Graceful Discard is set to On. The red frame percent limits the number of red frames the network is responsible to deliver.
PVC Loopback Status (Fwd/Rev)	Displays the current loopback state.
Reroute Balance	Choose *Enable* to allow this circuit to use reroute tuning. This feature enables the switch to redistribute PVCs across trunks based on OSPF updates and cost metrics. You must first configure the reroute tuning parameters for the selected switch. If you *Disable* this option, this circuit does not use the reroute tuning parameters.
Bandwidth Priority	Set a value from 0 through 3 where 0 is the default and indicates the highest priority.
Bumping Priority	Set a number from 0 through 7 where 0 is the default and indicates the highest priority.
QuickPath Segment Size (Bytes)	Not supported.

3. Choose OK to accept the circuit parameters and send the configuration file to the switch (provided the switch is communicating with the NMS). The Set All PVCs on Map dialog box reappears.

Defining an IP Flow Profile

To define an IP QoS PVC Flow Profile:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All QoS Profiles ⇒ Set All QoS Profiles. The Set All QoS Profiles dialog box appears (Figure 5-6).

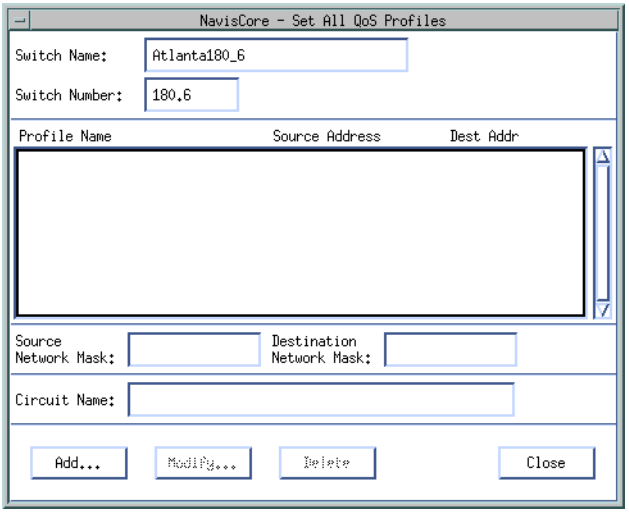

Figure 5-6. Set All QoS Profiles Dialog Box

The following table describes each of the Set All QoS Profiles buttons.

Button	Function
Add	Enables you to add an IP QoS PVC flow profile.
Modify	Enables you to modify an IP QoS PVC flow profile.
Delete	Enables you to delete an IP QoS PVC flow profile.

3. Choose the Add button. The Add IP QoS PVC Flow Profile dialog box appears (Figure 5-7).

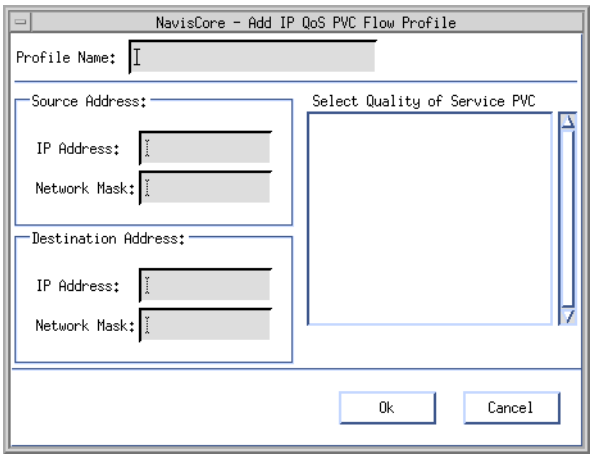

Figure 5-7. Add IP QoS PVC Flow Profile Dialog Box

4. Complete the fields as described in Table 5-6.

Table 5-6. Add IP QoS PVC Flow Profile Fields

Field	Action/Description
Profile Name	Enter a profile name that associates the profile with a logical port.
Source Address	
IP Address	Enter the IP address of the network or host that sends the packet.
Network Mask	Enter the network mask that applies to the source address.
Destination Address	
IP Address	Enter the IP address of the network or host that receives the packet.
Network Mask	Enter the network mask that applies to the destination address.
Select Quality of Service PVC	Select the PVC to which you want to assign the IP QoS PVC flow profile.

You do not have to specify both the destination and source address. However, you must specify at least one.

5. Choose OK.

6. At the Set All QoS Profiles dialog box, choose Close.

Assigning an IP Flow Profile to a Logical Port

> Before you assign an IP QoS PVC flow profile to the ingress IP logical port, you must enable QoS on the IP logical port. See Chapter 3, "Configuring IP Logical Ports and IP Servers" for information.

To assign an IP QoS PVC flow profile to a logical port:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, choose Ascend IP Parameters ⇒ Set All QoS Profiles ⇒ Set All Logical Port QoS Profiles. The Set All Logical Port QoS Profiles dialog box appears (Figure 5-8).

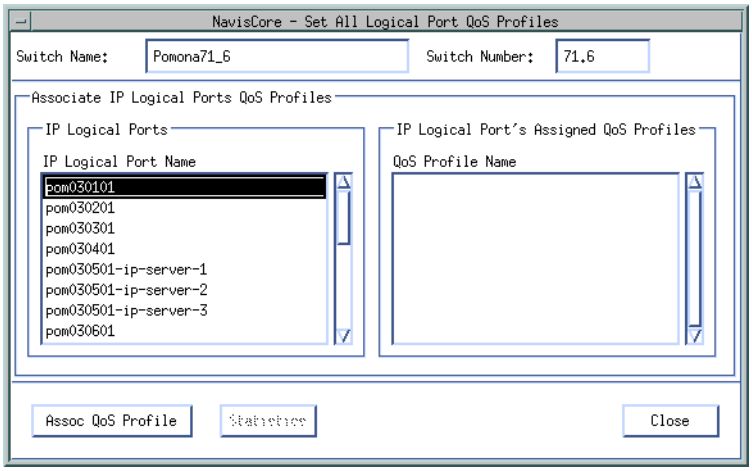

Figure 5-8. Set All Logical Port QoS Profiles Dialog Box

3. Select the IP logical port with which you want to associate an IP QoS PVC flow profile and choose the Assoc QoS Profile button. The Associate LPort QoS Profile dialog box appears (Figure 5-9).

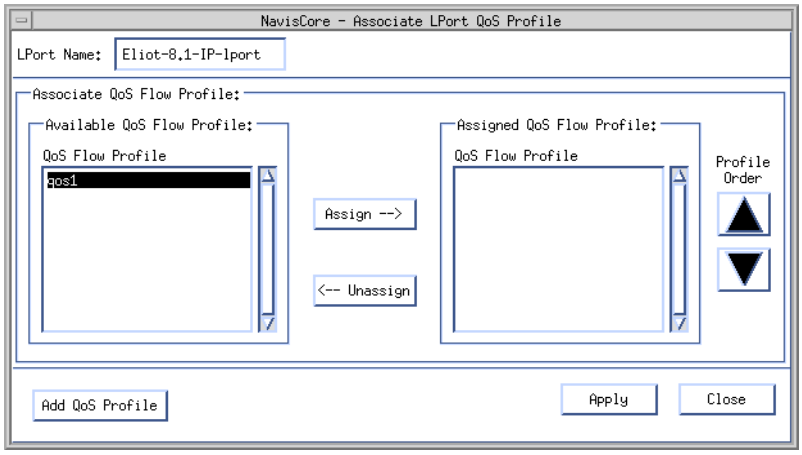

Figure 5-9. Associate Lport QoS Profile Dialog Box

4. Specify the parameters as described in Table 5-7.

Table 5-7. Lport QoS Profile Fields

Field	Action Description
LPort Name	Displays the name of the logical port.
Available QoS Flow Profile	Lists all current IP QoS PVC flow profiles that are available.
Assigned QoS Flow Profile	Lists all current IP QoS PVC flow profiles assigned to this logical port.
Assign Button	Enables you to assign an IP QoS PVC flow profile to the logical port.
Unassign Button	Enables you to delete an IP QoS PVC flow profile from a logical port.
Profile Order Buttons	Enables you to assign the flow profile's order of importance. In cases where the IP flow profile matches multiple profiles, the first profile is always used.
Add QoS Profile	Enables you to define additional IP QoS PVC flow profiles. See "Defining an IP Flow Profile" for information.

5. From the Available QoS Flow Profile list box, select the profile to associate with the logical port and choose Assign.

To delete a profile from the list, select a profile from the Assigned QoS Flow Profile list box and choose Unassign.

6. Choose Apply.

To create another IP QoS PVC Flow Profile, choose Add QoS Profile and see "Defining an IP Flow Profile" on page 5-13.

7. Choose Close.
8. At the Set All Logical Port QoS Profiles dialog box, choose Close.

6

Configuring Static ARP Entries

This chapter describes how to define Static ARP Entries. When you define Static ARP entries, you create a table that matches IP addresses to specific MAC, DLCI, or VPI-VCI addresses. The hardware address you define depends on the link type.

Address Resolution Protocol

A node requires the following information in order to communicate with another node:

- IP address of the destination node
- Hardware address of the destination node (DLCI for Frame Relay and VPI/VCI for ATM)

When an interface is configured for Ethernet, the IP addresses of the destination nodes are known (the hardware addresses are not known). When an interface is configured for Frame Relay, the hardware addresses of the destination nodes are known. IP services uses ARP and InARP to resolve the lack of a hardware or IP address.

Defining a Static ARP Entry

Before you create a static ARP entry, make sure that the hardware address that you plan to use as the Static ARP entry (either DLCI or VPI/VCI) has been already defined for the IP logical port on the Set IP Protocol Connection ID dialog box. See one of the following sections for details:

- "Setting the DLCI for Frame Relay Logical Ports" on page 3-14
- "Setting the VPI/VCI for ATM Logical Ports" on page 3-16

To define a static ARP entry:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Static ARP Entries. The Set All Static ARP Entries dialog box appears (Figure 6-1).

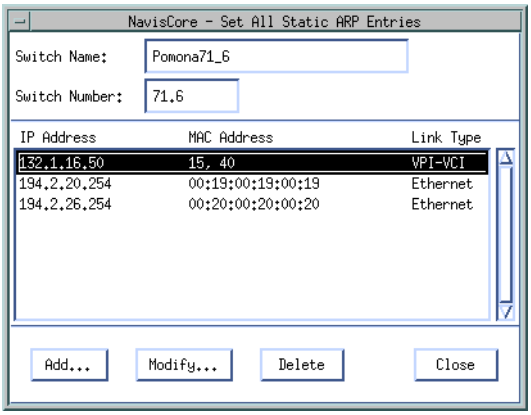

Figure 6-1. Set All Static ARP Entries List Dialog Box

The Set All Static ARP Entries dialog box displays the following buttons.

Button	Function
Add	Enables you to add a static ARP entry.
Modify	Enables you to modify a static ARP entry.
Delete	Enables you to delete a static ARP entry.

3. Choose the Add button. The Set Static ARP dialog box appears (Figure 6-2).

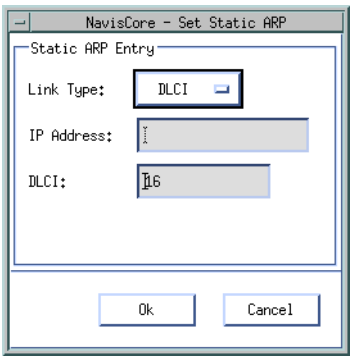

Figure 6-2. Set Static ARP Dialog Box

4. Select the link type and complete the fields described in Table 6-1.

The hardware address that you specify as the Static ARP entry (either DLCI or VPI/VCI) must have already been specified for the IP logical port on the Set IP Protocol Connection ID dialog box. See one of the following sections for details:

- "Setting the DLCI for Frame Relay Logical Ports" on page 3-14
- "Setting the VPI/VCI for ATM Logical Ports" on page 3-16

Table 6-1. Static ARP Fields

Link Type	Field	Action/Description
DLCI	IP Address	Enter the IP address of the neighbor.
	DLCI	Enter the DLCI used for the neighbor. Valid values range from 0 through 937. A DLCI is a 10-bit address that identifies PVCs.
VPI-VCI	IP Address	Enter the IP address of the neighbor.
	VPI	Enter the VPI used for the neighbor. Valid values range from 0 to 255. A VPI is an 8-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.
	VCI	Enter the VCI used for the neighbor. Valid values range from 0 to 255. A VCI is a 16-bit field in the ATM cell header that is used as an addressing identifer to route cell traffic.
Ethernet	IP Address	Enter the IP address of the neighbor.
	MAC Address	Enter the MAC Address used for the neighbor. A MAC address is a standardized data link layer address that is required for every port or device that connects to a LAN.

5. Choose OK. The ARP table entry is created for these addresses.
6. At the Set IP ARP List dialog box, choose Close.

7

Configuring RIP

This chapter describes how to configure Routing Information Protocol (RIP) parameters on an IP logical port. RIP is a distance vector protocol, which bases all routing decisions on the distance to the destination.

Configuring RIP at the Logical Port

To configure RIP at the logical port:

1. Add an IP logical port and interface. For more details on these procedures, see "Adding an IP Logical Port" on page 3-8.
2. Choose Add RIP from the Set IP Interface Addresses dialog box. The Add RIP Interface dialog box appears (Figure 7-1 on page 7-2).

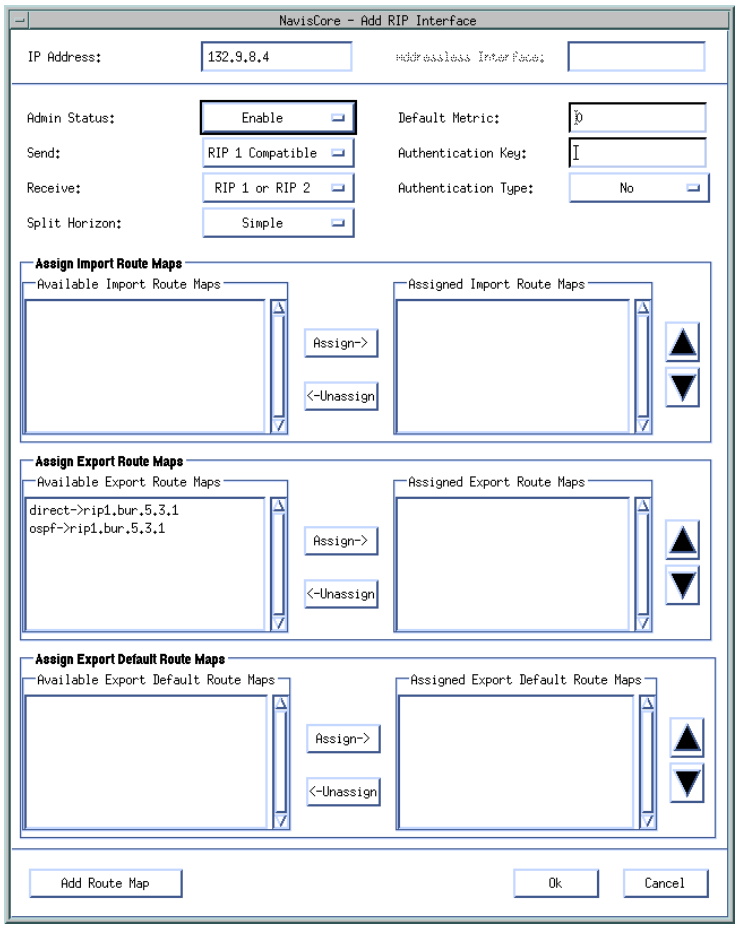

Figure 7-1. Add RIP Interface Dialog Box

3. To add a route map for this RIP interface, choose Add Route Map. See Chapter 11, “Configuring Route Maps” for more information on route maps.

4. To define the RIP interface, specify the field values as described in Table 7-1.

Table 7-1. RIP Interface Fields

Field	Action/Description
IP Address	The IP address for this interface.
Admin Status	Select one of the following options: *Enable* – Indicates that the port is activated for RIP and RIP packets can be exchanged over this logical port. *Disable* – Indicates that the port has not been activated for RIP or that the port is offline for diagnostics. An IP interface with an Admin Status of *Disable* cannot exchange RIP packets.
Send	Possible values are: *Disable, RIP 1, RIP 1 Compatible,* or *RIP 2.* RIP 1 Compatible is the default value.
Receive	Possible values are: *RIP 1, RIP2, RIP 1 or RIP 2,* or *Disable. RIP1 or RIP 2* is the default value.
Split Horizon	Split horizon is a method for avoiding common situations that require *counting to infinity.* Specify one of the following options: *Disable* – Indicates that split horizon will not be used. *Simple* – Indicates that split horizon will be used. The simple form of split horizon specifies that if a router learns of a route from an update received on the link, it does not advertise that route on updates that it transmits to the link. *Poison Reverse* – Is a stronger form of split horizon. In this form, routers do not omit destinations learned from an interface. Instead, they include these destinations, but advertise an infinite cost to reach them. This option increases the size of routing updates. In addition, it provides a positive indication that a specific location is not reachable through a router.
Default Metric	A variable that specifies the metric that is used for the default route entry in RIP updates that originate on this interface. A value of zero indicates that no default route should be originated.
Authentication Key	Do not specify this value if you specified a value of *No* as the authentication type. If you specified a value of *Simple* or *MD5* as the authentication type, you must specify the authentication password in this field.

Table 7-1. RIP Interface Fields (Continued)

Field	Action/Description
Authentication Type	This value specifies the type of authentication that RIP uses as a security measure to ensure that this logical port and router are exchanging information with proper neighbors. Possible values are *No*, *Simple*, or *MD5*. *No* – Specifies that no authentication will be performed. *Simple* – Specifies a simple password authentication method that enables you to designate a password that is part of all RIP messages on an interface-by-interface basis. When a router receives a message on an interface that is using simple password authentication, it checks the incoming RIP message to ensure that the proper password is included in the message. If the password is correct, the message is processed normally. If the password is not part of the incoming message or an incorrect password is used, the message is ignored and dropped. *MD5 Authentication* – Specifies the Message Digest Algorithm Version 5 (MD5) authentication. This method is similar to the simple password method, however, the password is never transmitted. Instead, the router uses the MD5 algorithm to create a message digest of the password. The message digest is sent instead of the password. This method prevents the password from being read during transmission.
Available Import Route Maps	The import route maps that are available for assignment to this RIP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Maps”. To display the parameters for any listed route map, double-click on the map.
Assigned Import Route Maps	The import route maps that are assigned to this RIP interface. All incoming routes on this RIP interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Available Export Route Maps	The export route maps that are available for assignment to this RIP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Maps”. To display the parameters for any listed route map, double-click on the map.

Table 7-1. RIP Interface Fields (Continued)

Field	Action/Description
Assigned Export Route Maps	The export route maps that are assigned to this RIP interface. All outgoing routes on this RIP interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

5. Choose OK. NavisCore adds the RIP interface and returns to the Set IP Interface Addresses dialog box.

8

Configuring BGP Parameters

This chapter provides an overview of the Border Gateway Protocol (BGP) and describes how to perform the following tasks:

- Defining BGP switch parameters
- Defining BGP neighbors
- Defining BGP aggregates
- Setting IP Loopback addresses

About BGP

BGP is a protocol that exchanges routing information between autonomous systems, which is a set of routers having a single routing policy running under a single technical administration. BGP advertises routes between external BGP neighbors or peers, unlike Interior Gateway Protocol (IGP), which advertises routes between internal peers within the same autonomous system, such as Open Shortest Path First (OSPF) and Routing Information Protocol (RIP).

See Figure 8-1 for an example of AS relationships.

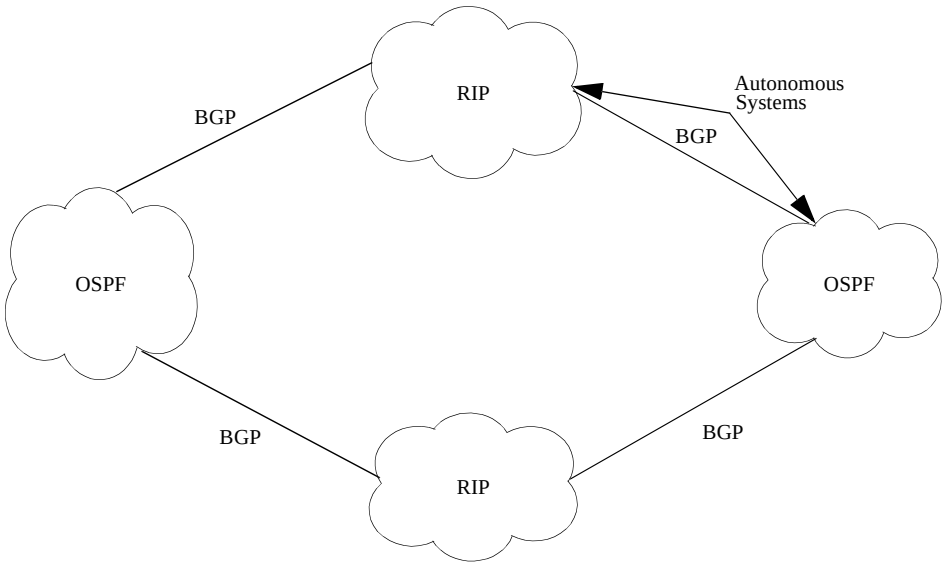

Figure 8-1. Autonomous System Examples

BGP Peers and Route Updates

BGP is considered a path-vector protocol because it carries a sequence of autonomous-system numbers that indicate the path a *route* has taken. When you define an autonomous system, you specify the networks to which a BGP peer sends route information. The network administrator uses a set of BGP parameters as tie breakers to indicate the routes that BGP should select as the best path.

BGP peers form a connection between each other, exchanging messages to open and confirm a TCP-based connection. Peers exchange route-update messages, which contain network reachability, path attributes, and preferred-route information. If there is disagreement between the peers, BGP sends an error to each peer and the connection is not established.

BGP updates route changes in a routing table. If routing information changes, BGP informs the peers by removing invalid routes and adding the new route information. If no changes occur, BGP peers exchange keep-alive messages to ensure the connection is alive.

Configuring IBGP

Typically, OSPF and RIP are used as the interior gateway protocol within the autonomous system. However, you can use BGP as the IGP. You can configure Interior Border Gateway Protocol (IBGP) the following ways:

- Full Mesh IBGP
- Route Reflection

Full Mesh IBGP

In a full mesh IBGP, all IBGP neighbors within an autonomous system must be connected to exchange route update information. However, this is not the preferred configuration due to limited computing resources in a switch environment.

Figure 8-2 displays a full mesh IBGP.

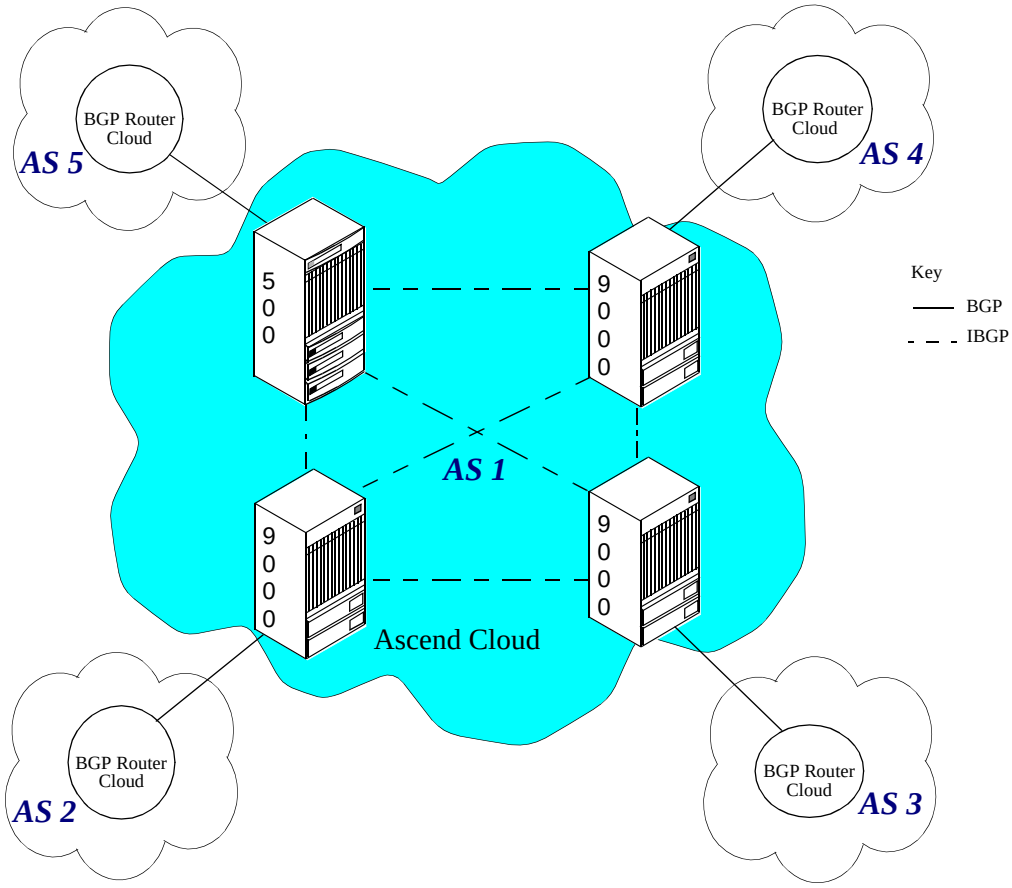

Figure 8-2. Full Mesh Interior Border Gateway Protocol Example

Route Reflection

Route reflection is a better alternative to full mesh IBGP. In route reflection, a BGP switch is designated as the route reflector, sending or *reflecting* received route information to all internal neighbors (or peers). There are two groups of route reflection peers:

- Client peers
- Non-client peers

When comparing the two groups, client peers do not have to be meshed, while non-client peers must be fully meshed together. Client peers are grouped into a *cluster* and communicate with each other. Client peers cannot communicate with non-client peers (peers outside of their cluster) but must communicate with the route reflector that belongs to the non-client peers' cluster.

Figure 8-3 illustrates an example of a route reflection configuration.

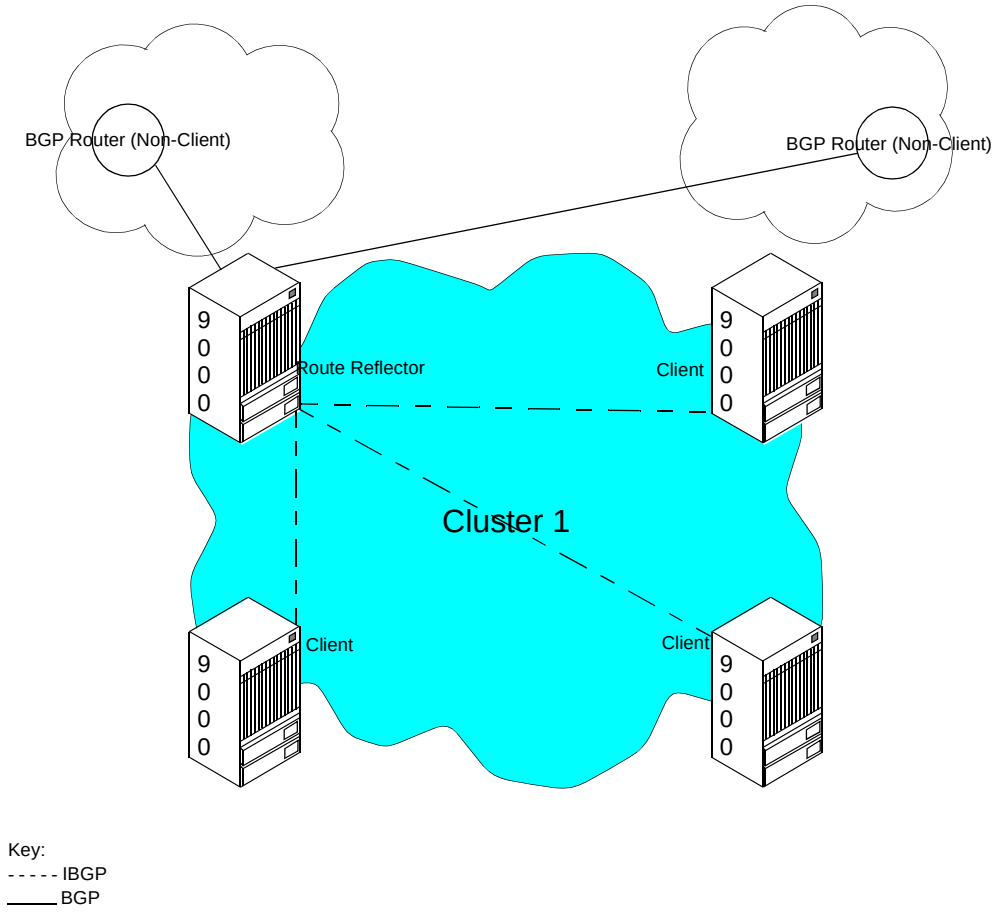

Figure 8-3. Route Reflection Example

For every route update received from an advertiser peer, the route reflector does one of the following (provided the best path selection is applied first):

- If the advertiser peer is a non-client, then the route reflector reflects the route to all non-clients.
- If the advertiser peer is a client peer, then the route reflector reflects the route to all non-client peers and all client peers other than the original advertiser.
- If the advertiser peer is an external BGP peer, then the route reflector reflects the route to all clients and non-clients (normal BGP operation).

Route reflection defines the following attributes for detection and avoidance of path loops:

ORIGINATOR_ID — This attribute is the router ID of the route originator in the local AS.

CLUSTER_LIST — This attribute is a sequence of cluster ID values that represent the reflection path the route passed.

Autonomous systems may have multiple route reflectors. Route reflectors communicating with each other are considered non-client peers and should be fully meshed.

BGP Aggregates

An aggregate is the combining of specific network addresses to less specific ones. This reduces the size of the routing table. Aggregates do the following:

- Reduce the size of the BGP routing table
- Provide better network control over network instability
- Provide a better mechanism to maintain route updates across areas

Aggregate networks are classless (CIDR) and configured with a network prefix and mask. During the route update process, BGP scans the entire routing table for networks that are part of the configured aggregate network. If matches are found, BGP forms the aggregate networks and advertises aggregate routes to peers.

Configuring BGP

This section describes how to set BGP parameters on the switch level and includes the following tasks:

- Define BGP switch parameters
- Define a BGP neighbor and assign route maps
- Define a BGP aggregate

Defining BGP Switch Parameters

To define BGP switch parameters:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All BGP ⇒ Set All BGP Parameters. The Set BGP dialog box appears (Figure 8-4).

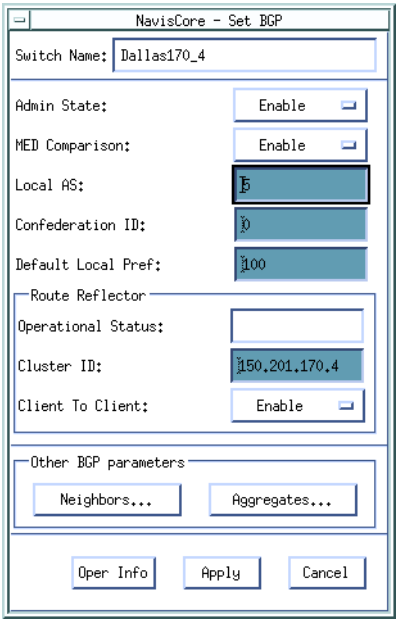

Figure 8-4. Set BGP Dialog Box

The Set BGP dialog box provides the following option buttons:

Button	Function
Neighbors...	Choose this button to define BGP neighbor parameters. For more information, see "Defining a BGP Neighbor and Assigning a Route Map" on page 8-8.
Aggregates...	Choose this button to define BGP aggregate parameters. For more information, see "Defining a BGP Aggregate" on page 8-14.

3. Complete the fields described in Table 8-1.

Table 8-1. BGP Parameter Fields

Field	Action/Description
Admin State	Select one of the following options: *Enable* – Allows the selected switch to exchange route updates using BGP. *Disable* – Prevents the selected switch from exchanging route updates using BGP.
MED Comparison	Select one of the following options: *Enable* – Allows you to use a multi-exit discriminator (MED) in the route selection process. MED allows BGP to communicate preferred path information to external neighbors when the autonomous system has multiple exits to another autonomous system. *Disable* – Prevents the use of MED in the route selection process.
Local AS	Enter a value between 1 and 65535. This parameter is the switch's autonomous system number.
Confederation ID	Confederation provides a method for dividing a larger AS into multiple smaller sub-ASs. If you are using BGP confederation, this value enables you to specify a unique confederation identifier. The ID may be any unique number within the range of 1 to 65535.
Default Local Pref	Enter a value between 1 and 4294967295 (the default is 100). This value is sent to internal neighbors. A local preference allows you to rank a route according to its importance. The local preference is compared to other routes that have the same destination. A higher local preference indicates the route is preferred.

Table 8-1. BGP Parameter Fields (Continued)

Field	Action/Description
Route Reflector	
Operational Status (READ ONLY)	This parameter identifies whether or not this peer is a route reflector. If it is, the peer forwards route information to all clients. The route reflector is implicitly defined when you define any of its peers to be a route reflector client.
Cluster ID	Enter the internal IP address of the selected switch if the switch is a route reflector in a cluster, which contains more than one route reflector. A cluster is a group of client peers that communicate with a BGP route reflector. A cluster ID specifies the cluster.
Client To Client (For IBGP peers only)	Select one of the following options: *Enable* – If you enable this parameter, any routes that are received by the selected switch from a client will be sent to all other clients. Enable is the default. *Disable* – If you disable this parameter, any routes that are received by the selected switch from a client will not be sent to all other clients. ***Note****: Disable this parameter if all clients are fully meshed.*

4. Choose Oper Info to display a status message about the selected port.
5. Choose OK.

Defining a BGP Neighbor and Assigning a Route Map

In addition to defining BGP neighbors, you must assign route filters to these BGP nodes. Route maps control the flow of route updates. You use a route filter to selectively accept, reject, advertise, or hide routes. See Chapter 11, "Configuring Route Maps" for details on defining route maps.

To define a BGP neighbor to a switch:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All BGP ⇒ Set All BGP Neighbors. The Set All BGP Neighbors dialog box appears (Figure 8-5).

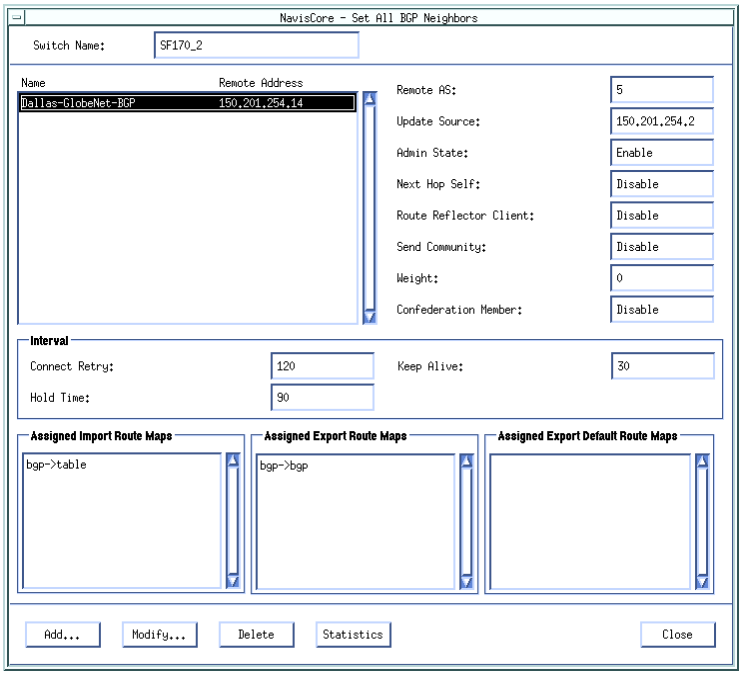

To display the parameters for a listed route map, double-click on the route map.

Figure 8-5. Set All BGP Neighbors Dialog Box

The Set All BGP Neighbors dialog box displays the following buttons:

Button	Function
Add	Enables you to add a BGP neighbor.
Modify	Enables you to modify a BGP neighbor.
Delete	Enables you to delete a BGP neighbor.
Statistics	Use the Statistics option to display BGP peer connection statistics. For more information, see the *NavisCore Diagnostic and Troubleshooting Guide.*

3. Choose Add.

If you selected a switch that has an AS of zero, the following error message appears:

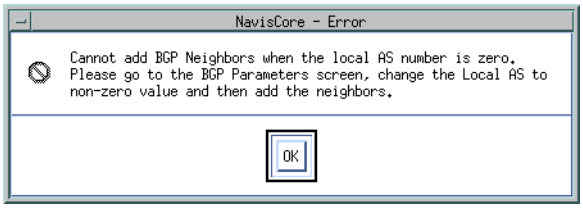

Figure 8-6. BGP Neighbor Error Message

Go to "Defining BGP Switch Parameters" on page 8-6 and change the switch's AS number to a non-zero value.

If you selected a switch that has a non-zero AS value, the Add BGP Neighbor dialog box appears (Figure 8-7):

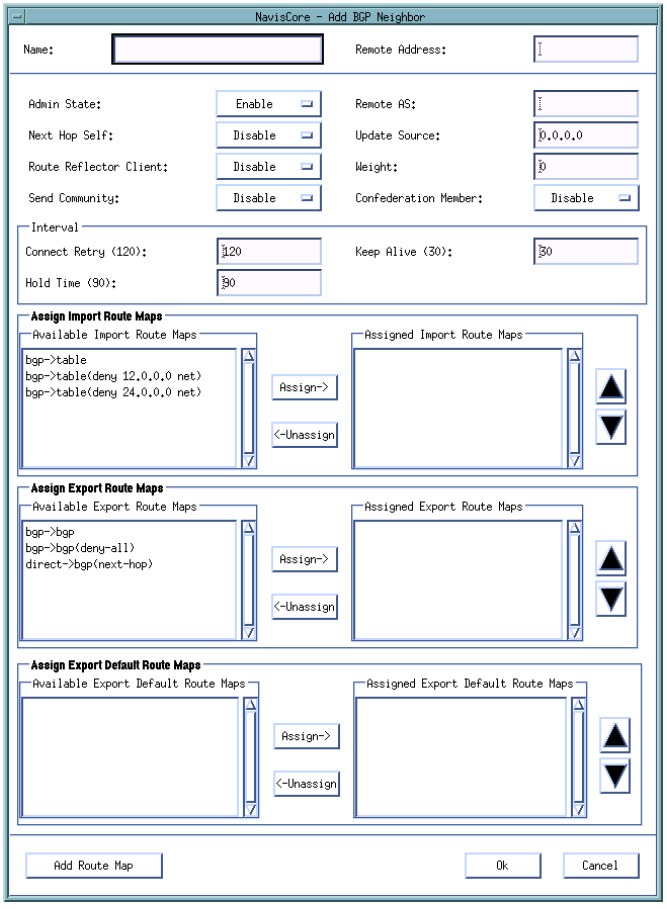

Use the arrow buttons to specify the sequence that IP Navigator uses to filter routing information.

Figure 8-7. Add BGP Neighbor Dialog Box

4. Specify the values as listed in Table 8-2.

Table 8-2. BGP Neighbor Fields

Field	Action/Description
Name	Enter the name of the BGP neighbor.
Remote Address	Enter the IP address of the BGP neighbor.
Admin State	Select one of the following options: *Enable* – Activates the connection between the selected switch and this BGP neighbor. *Disable* – Deactivates the connection between the selected switch and this BGP neighbor.
Remote AS	Enter a value between 1 and 65535. This value is the neighbor's remote AS number.
Next Hop Self	Select one of the following options: *Enable* – For IBGP peers, enabling this parameter forces BGP to advertise the local address of the BGP connection as the next hop. For EBGP peers, BGP always advertises the local address as the next hop; therefore you do not need to enable next hop self for EBGP peers. *Disable* – Disabling this parameter allows BGP to determine the next hop.
Update Source	Enter a valid IP address for the Update Source address, which is the source address for the BGP TCP connection.
Route Reflector Client	Select one of the following options: *Enable* – If you enable this parameter, the selected switch's neighbor is defined as a route reflector client, implicitly making the selected switch a route reflector. *Disable* – If you disable this parameter, the selected switch's neighbor is not defined as a route reflector client. In addition, if you disable this parameter on all of the selected switch's BGP neighbors, the selected switch is not defined as a route reflector. However, if the route reflector client is enabled on at least one BGP neighbor, the selected switch is still considered a route reflector.
Weight	Enter a value between 0 and 65535 (the default value is zero).This parameter represents the path weight (received by the neighbor) that is applied to every EBGP route. IP Navigator applies the weight value to EBGP routes only. It does not use the weight value for IBGP routes.

Table 8-2. BGP Neighbor Fields (Continued)

Field	Action/Description
Confederation Member	Select one of the following options: *Enable* – If you enable this parameter, you specify that this BGP neighbor is a confederation member. Confederation provides a method for dividing a large AS into multiple, smaller sub-ASs. While BGP confederations provide some advantages, the migration from a nonconfederation to a confederation design requires reconfiguration of the routers and a major change in the local topology. *Disable* – If you disable this parameter, you specify that this BGP neighbor is not part of a confederation.
Send Community	Select one of the following options: *Enable* – Enables you to send community attributes of all updates to this neighbor. A community is a group of destinations that share some common property. A community is not restricted to one network or autonomous system; it has no physical boundaries. You use community attributes to simplify routing policies by identifying routes based on the logical property rather than IP prefix or AS number. *Disable* – Disables the sending of community attributes of all updates to this neighbor.
Interval The default value for each interval field in the Add BGP Neighbor dialog box (Figure 8-7) is in parentheses.	
Connect Retry(120)	Enter a value between 1 and 65535 (the default is 120). This parameter is the time, in seconds, that BGP waits before it tries to connect to this neighbor. The number of connection retries due to errors are generated with no regard to this value. The initial value is 60 seconds, which is doubled for each retry after that.
Keep Alive(30)	Enter a value between 0 and 21845 (the default is 30). This parameter is the time, in seconds, between consecutive keep alive messages sent to this neighbor. This event occurs after a connection is established. Keep alive messages are sent periodically between BGP neighbors to ensure that the connection is still alive.
Hold Time(90)	Enter either a value of 0, or a range of 3 to 65535 (the default is 90). The value 0 indicates not to use hold time with this neighbor. This parameter represents the time, in seconds, BGP holds before considering the connection to be down if messages are not received from this neighbor.
Assign Import Route Maps	
Available Import Route Maps	The import route maps that are available for assignment to this BGP neighbor. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Maps”. To display the parameters for any listed route map, double-click on the map.

Table 8-2. BGP Neighbor Fields (Continued)

Field	Action/Description
Assigned Import Route Maps	The import route maps that are assigned to this BGP neighbor. All incoming routes on this BGP neighbor are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Assign Export Route Map	
Available Export Route Maps	The export route maps that are available for assignment to this BGP neighbor. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Maps". To display the parameters for any listed route map, double-click on the map.
Assigned Export Route Maps	The export route maps that are assigned to this BGP neighbor. All outgoing routes on this BGP neighbor are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Assign Export Default Route Maps	
Available Export Default Route Maps	The export default route maps that are available for assignment to this BGP neighbor. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Maps". To display the parameters for any listed route map, double-click on the map.
Assigned Export Default Route Maps	The export default route maps that are assigned to this BGP neighbor. All outgoing routes on this BGP neighbor are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps s*hould be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

5. In the Available Import Route Maps List box, specify the import route map and choose assign.

   To remove an import route map from the list, select the import route map from the Assigned Import Route Maps List box and choose Unassign.

6. In the Available Export Route Maps List box, specify the export route map and choose assign.

   To remove an export route map from the list, select the export route map from the Assigned Export Route Maps List box and choose Unassign.

7. In the Available Export Default Route Maps List box, specify the export default route map and choose assign.

   To remove an export default route map from the list, select the export default route map from the Assigned Export Default Route Maps list box and choose Unassign.

8. Choose OK.

9. To add an additional network access list, choose Add Route Map. See "Adding Route Maps" on page 11-16 for more information.

10. In the Set All BGP Neighbors dialog box, choose Close.

Defining a BGP Aggregate

To define BGP Aggregates:

1. From the network map, select the appropriate switch icon.

2. From the Administer menu, choose Ascend IP Parameters ⇒ Set All BGP ⇒ Set All BGP Aggregates. The Set All BGP Aggregates dialog box appears (Figure 8-8).

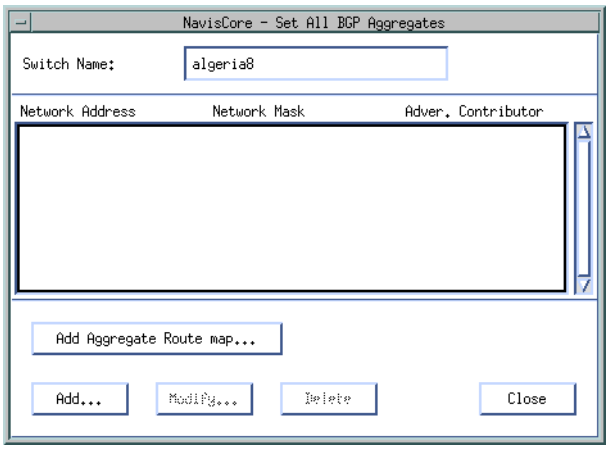

Figure 8-8. Set All BGP Aggregates Dialog Box

The Set All BGP Aggregates dialog box displays the following buttons.

Button	Function
Add Aggregate Route Map	Enables you to add an Aggregate Route Map. For more information, see Table 11-26 on page 11-48.
Add	Enables you to add a BGP aggregate.
Modify	Enables you to modify a BGP aggregate.
Delete	Enables you to delete a BGP aggregate.

3. At the Set All BGP Aggregates dialog box, choose Add. The Add BGP Aggregates dialog box appears (Figure 8-9).

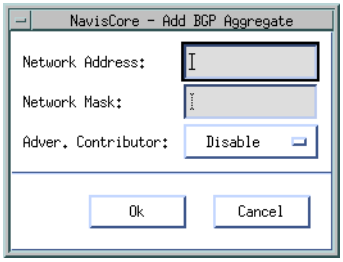

Figure 8-9. Add BGP Aggregate Dialog Box

4. Specify the values described in Table 8-3.

Table 8-3. BGP Aggregate Fields

Field	Action/Description
Network Address	This parameter is the aggregate network IP address.
Network Mask	This parameter is the aggregate network mask.
Adver. Contributor	Select one of the following options: *Enable* – Enabling this parameter allows you to advertise components of the aggregate network. *Disable* – Disabling this parameter enables you to stop advertising components of the aggregate network.

5. Choose OK.
6. At the Set All BGP Aggregates dialog box, choose Close.

Setting IP Loopback Addresses

The Set IP Loopback Address function enables you to establish an IP loopback address that is not associated with any physical port. Because the loopback address is independent of a physical interface, the status of the physical link does not affect the IP loopback address. If you use an IP loopback address as a BGP neighbor address, you ensure that the BGP connection will not go down.

To set an IP loopback address:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, choose Ascend IP Parameters ⇒ Set All IP Loopback Addresses. The Set All IP Loopback Addresses dialog box appears (Figure 8-10).

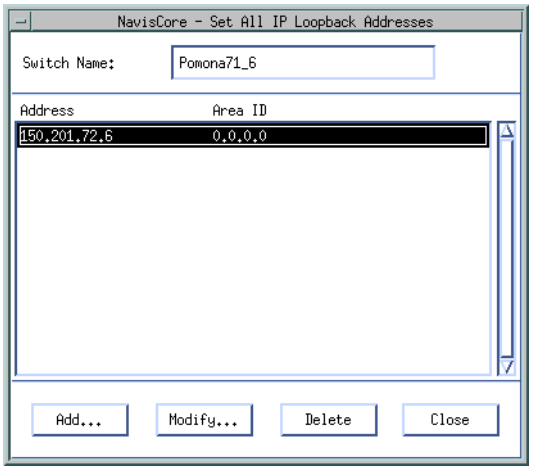

Figure 8-10. Set All IP Loopback Addresses Dialog Box

3. Choose Add. The Add IP Loopback Address dialog box appears (Figure 8-11).

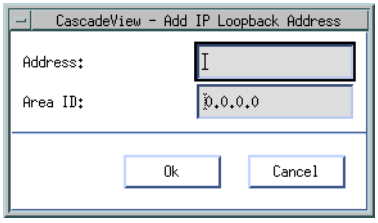

Figure 8-11. Add IP Loopback Address Dialog Box

4. Enter the IP address.
5. Enter the Area ID.
6. Choose OK. The Set All IP Loopback Addresses dialog box reappears and the new IP loopback address is included in the list.

9

Configuring OSPF Parameters

This chapter provides an overview of the Open Shortest Path First (OSPF) protocol and describes how to perform the following tasks:

- Configuring OSPF at the Logical Port
- Configuring OSPF Parameters at the Switch
- Configuring IP Parameters
- Defining an OSPF Neighbor
- Defining an OSPF Area Aggregate
- Defining an OSPF Virtual Link
- Configuring an OSPF Route Map
- Configuring Multiple OSPF Areas

About OSPF

The OSPF protocol is a link-state routing protocol. With link-state routing protocols, routers maintain a link-state database that contains topology current link-state information. This information enables routers to determine the best routes to each destination network in the autonomous system. The OSPF protocol has the following advantages over the Routing Information protocol (RIP):

Authentication — Provides security. Only an authorized router can generate route updates to other routers.

Type of Service (TOS) — Enables your network to make routing decisions based on the quality of service required by a host application.

Areas — Restricts flooding to configured areas, thereby reducing the database size.

The Link-State Database

OSPF *floods* routers with link-state advertisements (LSAs) which contain topology information. Routers store LSAs in the link-state databases. Flooding ensures that all routers have identical databases and the same topology information.

Link-state databases include:

- Known router addresses
- Known links and their associated costs
- Known network addresses

Routers use the link-state database and Dijkstra's algorithm (algorithm used to calculate best routes) to determine the best route.

Designated Routers and OSPF Relationships

Designated routers are responsible for sending copies of the link-state database to routers in the network. When new routers send hello packets to the designated router, the designated router responds with an acknowledgment message. The new router then sends a database description packet requesting a copy of the link-state database. The designated router responds by sending a database description packet that contains a copy of the link-state database to the new router.

In addition, designated routers:

- Monitor the health of adjacent routers
- Establish adjacencies

A backup designated router is defined in case the designated router goes down. The backup designated router keeps track of the same information as the designated router, but keeps silent. If the backup detects a failure of the designated router, it immediately becomes active.

OSPF Flooding Controls

Flooding is a reliable way to send link-state advertisements because many copies of the message travel through the network, ensuring that one message will arrive safely at each node. However, flooding causes significant network traffic. To reduce network traffic, OSPF implements the following flooding controls:

- The designated router is the only router that can generate link-state updates. This control reduces the number of copies created.
- Before forwarding OSPF link-state updates, the designated router checks its own link-state database to see if the update was received. If it was, the copy is discarded.
- OSPF supports areas where flooding is restricted. Smaller areas mean fewer copies of a message and less traffic.

Despite these benefits, flooding controls reduce the reliability of flooding. Flooding is reliable because many copies of the message travel through the network, ensuring a high probability that one message will arrive safely.

OSPF Areas

As networks grow large, the link-state database grows large as well. This causes problems for the following reasons:

- Increased memory space is consumed
- Route table generation becomes more processor-intensive
- It takes longer to:
  - Calculate link costs for more links
  - Calculate the spanning tree for a large network
  - Generate large routing tables required by large networks

To address large link-state databases, OSPF uses *areas*. An area is a group of OSPF routers that exchanges topology information. Designated routers only send link-state advertisements (LSAs) to routers that are part of the same area. If an autonomous system has one area, all routers in the autonomous system receive LSAs. However, if the autonomous system is divided into many areas, LSAs only go to the appropriate areas, thereby minimizing traffic and the link-state database size. The autonomous system works like a collection of smaller networks.

Because of flooding controls, the topology of one area is unknown by routers in another area. This means a router knows nothing of network topology outside its own area. Each area has a unique link-state database, and all routers in the given area should have the same database.

Figure 9-1 illustrates the concept of areas.

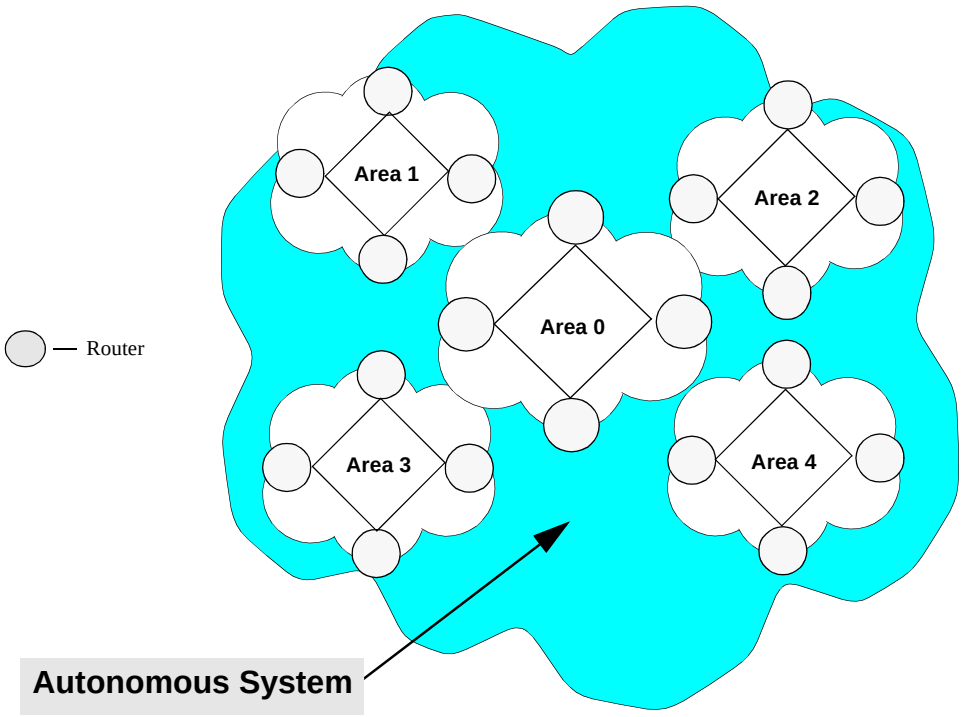

Figure 9-1. OSPF Areas

In Figure 9-1, the autonomous system is divided into five areas. Each area represents smaller networks within the autonomous system, and maintains separate link-state databases. Area 0 is the backbone and connects all areas within the autonomous system.

Area Aggregates

Area aggregates consolidate multiple routes (or addresses) within an area (or areas) into one single link state advertisement (LSA). This consolidation enables one advertisement representing a range of addresses within an area (or areas) to be broadcast.

Area aggregates:

- Reduce the size of the OSPF routing table
- Provide better control over network instabilities
- Provide a better mechanism to summarize route updates across areas
- Reduce memory requirements for link-state databases
- Reduce the cost of route calculation

- Have a maximum area size of 400 switches and routers, or 1,000 interfaces

The Backbone

The *backbone* is itself an area and is designated as area 0. OSPF requires the backbone to be contiguous to all areas in the autonomous system (AS). All other non-backbone areas (areas other than area 0) must have a connection to the backbone area.

Area Border Routers and Switches

Routers that belong to more than one area are referred to as *Area Border Routers* (ABRs). ABRs maintain separate link-state databases of each area to which they belong. A switch that spans one or more OSPF areas is considered to be an *area border switch* (ABS).

Virtual Links

OSPF requires that all OSPF areas be directly connected to the OSPF backbone area. However, you can use virtual links to logically connect physically separate portions of a network to the backbone.

OSPF uses virtual links for the following purposes:

- To connect areas that are not physically connected to the backbone.
- To patch the backbone if there is a break in backbone continuity.

The two endpoints of a virtual link are ABRs or ABSs. Figure 9-2 illustrates a network that connects Area 0.0.0.4 to the backbone through Area 0.0.0.3 by using a virtual link from Switch 7 to Switch 4.

About Clustering

Clustering is a way of grouping OSPF areas into subareas. Clustering enables you to use set increments (allows you to use a set of three bits of the internal IP address to assign a cluster address between 000 and 111, or 0 and 7) of the host ID address in different OSPF areas, while performing route aggregation at the ABS or ABR. A cluster forms a subset of an OSPF area. A cluster enables additional address aggregation at the ABS and reduces the size of the IP routing table, link-state database, and the number of summary link state advertisements (LSAs).

Use clustering only if you plan to do the following:

- Implement OSPF areas using switch software Version 5.0 or greater.
- Deploy new nodes with the same subnet addresses into multiple OSPF areas (for example, due to a lack of IP addresses)

In Version 2.3 and later versions of the NMS, you can define an IP address subnet as part of a cluster, define a cluster ID, and designate a switch as part of a cluster at switch deployment.

You assign a cluster ID to the IP address to be clustered. The cluster ID specifies the upper three bits of the host ID. As switches are added in that cluster ID the switch number/host ID in the IP address increments according to the cluster ID. For example, Table 9-1 shows the cluster ID IP-address range using 107.109.50.x as the default IP address.

Table 9-1. Cluster ID and IP Addresses

Cluster ID	IP Address Range
0	107.109.50.1 - 107.109.50.30
1	107.109.50.33 - 107.109.50.62
2	107.109.50.65 - 107.109.50.94
3	107.109.50.97 - 107.109.50.126
4	107.109.50.129 - 107.109.50.158
5	107.109.50.161 - 107.109.50.190
6	107.109.50.193 - 107.109.50.222
7	107.109.50.225 - 107.109.50.254

Figure 9-2 illustrates the benefits of clustering.

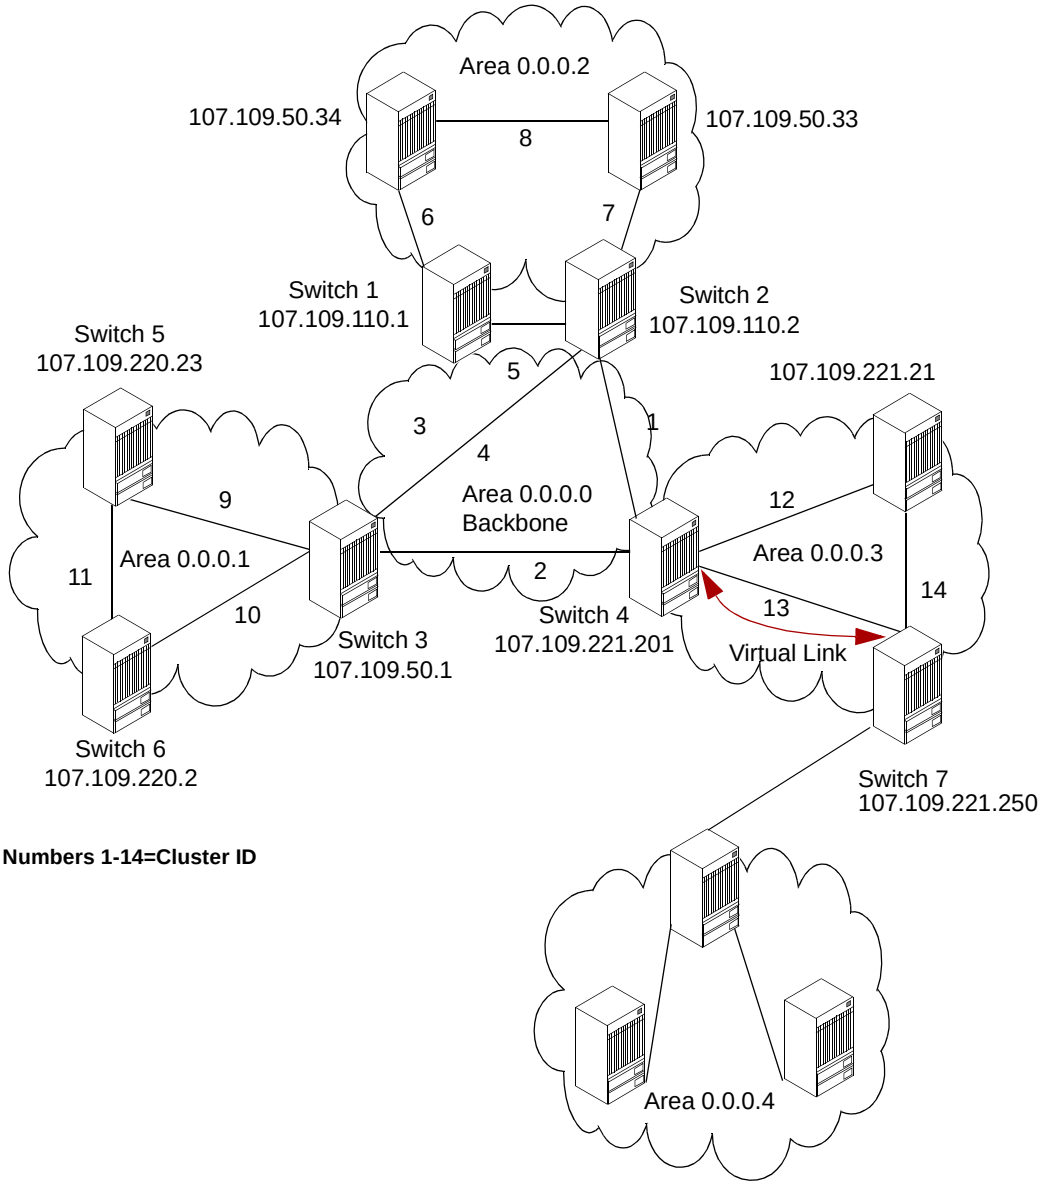

Figure 9-2. OSPF Area Configuration Example

Figure 9-2 shows a single OSPF backbone area with an Area ID of 0.0.0.0. All three non-backbone areas (0.0.0.1, 0.0.0.2, and 0.0.0.3) are directly connected to the OSPF 0.0.0.0 backbone area. Area 0.0.0.4 is connected to the backbone through the virtual link that is configured in Area 0.0.0.3.

Packets are forwarded between areas, from the source area, through the backbone, and then into the destination area.Ascend's OSPF area implementation assigns each trunk to a specific area. This provides maximum flexibility in setting area boundaries and changing area boundaries in the future.

Summary LSAs

IP addressing information is advertised across area boundaries in OSPF summary LSAs. Each summary LSA advertises a single range of IP addresses. The IP address ranges are configured in the ABSs.

For example, Area 1 is assigned a subnet of 107.109.220.0/*24*. The number 24 specifies a subnet mask of 24, so all IP addresses in the range 107.109.220.1-254 are sent as a single OSPF summary LSA. In this example, an OSPF summary LSA is sent from Area 2 for each address 107.109.50.33 and 107.109.50.34 without clustering. With clustering, the Areas 107.109.110.1 and 107.109.110.2 are configured for address range 107.109.50.32/27. A single summary LSA is sent for the 107.109.50.32/27 address range.

Address aggregation is not required, but when used, it results in fewer summary-LSAs. Fewer summary LSAs reduce the size of the routing table and OSPF link-state databases. In Figure 9-2, the Area 1 routing table would be 107.109.110.0/24, 107.109.50.33, and 107.109.50.34 without OSPF areas. If 107.109.50.0 is designated as cluster 1, the routing tables would have entries for 107.109.110.0/24 and 107.109.50.32/27.

If you anticipate a lack of network IP addresses and the use of a particular subnet address in multiple OSPF areas, you should add new switches to a cluster. Otherwise, you may not need to cluster.

OSPF Routing and Router Classifications

There are two types of routing:

- Intra-area routing
- Inter-area routing

Intra-area routing is routing within an area, and inter-area routing is routing between areas. These types of routing are performed by different classifications of routers, including:

Internal routers — Routers that are directly connected and belong to the same area. In addition, routers with interfaces connected only to the backbone are classified as internal routers.

Area border routers — Routers with links to more than one area, or between an area and backbone.

Backbone routers — Routers with an interface to the backbone. A backbone is either an area border router or an internal router.

AS boundary routers — Routers that connect an OSPF autonomous system to a region that uses a different routing protocol. AS boundary routers may be internal routers, area border routers, or backbone routers.

Figure 9-3 shows an example of OSPF routing and router classifications.

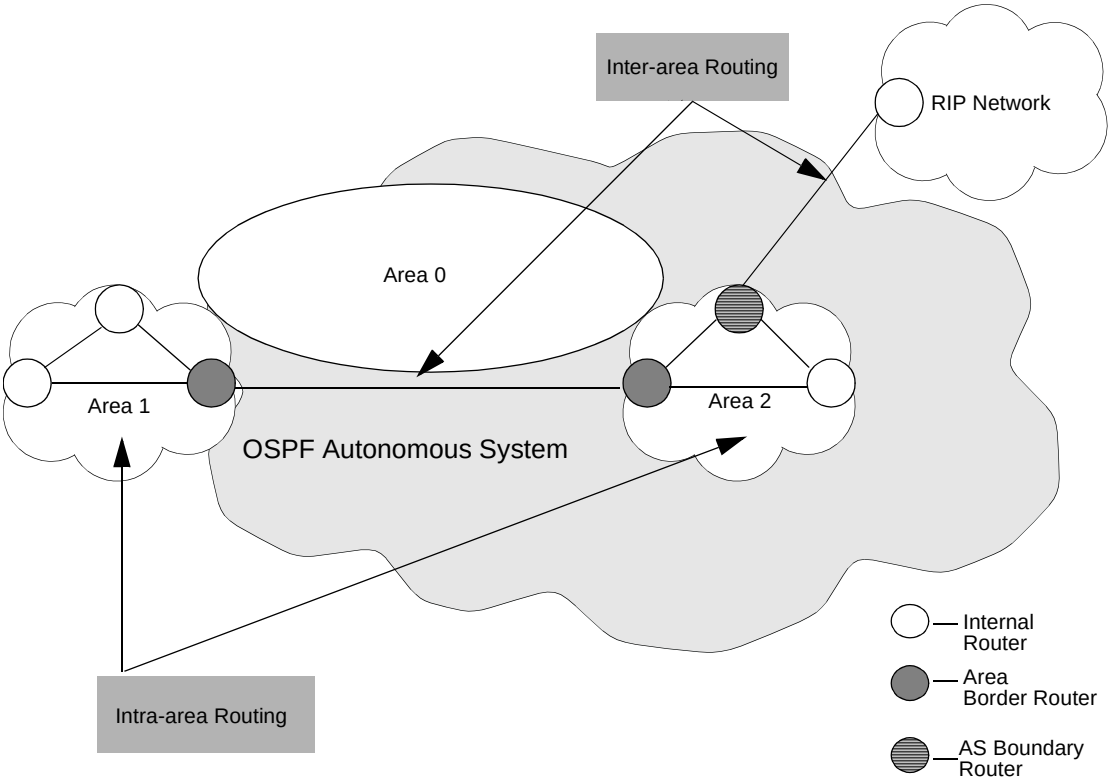

Figure 9-3. Router Classifications

Configuring OSPF

This section describes how to set OSPF parameters and includes the following tasks:

- Configuring OSPF parameters at the logical port
- Configuring OSPF parameters at the switch, including
  - Configuring IP parameters
  - Defining an OSPF neighbor
  - Defining an OSPF area aggregate
  - Defining an OSPF virtual link
  - Configuring OSPF route maps

Configuring OSPF at the Logical Port

To configure OSPF on the logical port:

1. Enable the logical port for IP services as described in "Configuring IP Logical Ports" on page 3-4.
2. Choose Add OSPF from the Set IP Interface Addresses dialog box. The Add OSPF Interface dialog box appears (Figure 9-4).

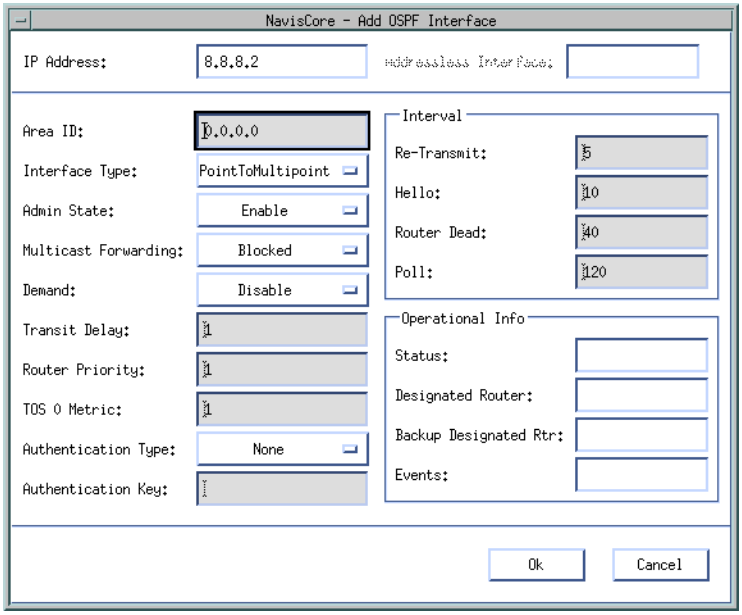

Figure 9-4. Add OSPF Interface

3. Complete the fields described in Table 9-2.

Table 9-2. Add OSPF Interface Fields

Field	Action/Description
IP Address	Displays the name assigned to the IP unicast address, with which this IP interface will communicate.
Addressless Interface	Enter the addressless interface. If the interface has an IP address, the value is 0.0.0.0. If the interface is addressless, the value is the logical port number or interface number.
Area ID	Enter the area ID (x.x.x.x) for the area in which you want to locate this interface. Area 0.0.0.0 is the network backbone area. Areas are collections of networks, hosts, and routers. The area ID identifies the area. ***Note***: *Area 1 is reserved for Ascend switches. If you configure the OSPF interface in Area 1, see "Configuring IP Parameters" on page 9-16.*
Interface Type	Select one of the following options: *Broadcast* – A broadcast network supports many routers and has a designated router that addresses a single physical message to all attached routers. The hello protocol dynamically discovers neighboring routers on these networks. *NBMA* – A non-broadcast multi-access (NBMA) network supports many routers, but does not have broadcast capability. This type of network requires full-mesh connectivity. *Point-to-Point* – A point-to-point network joins two routers together. The IP address of the neighboring routers interface is advertised. Hello packets are sent to the neighbor at regular intervals based on the value that you specify for the *Hello Interval* parameter. For more information on Hello Interval, see page 9-13. *Point-to-Multipoint* – A point-to-multipoint network supports multiple router connections, which are treated like point-to-point connections. The IP addresses of the remote routers interfaces are advertised. *Virtual Link* – A virtual-link network links areas that are not physically connected to the backbone and patches the backbone if a disconnect occurs in the backbone.
Admin State	Select one of the following options: *Enable* – This parameter allows this interface to communicate OSPF. In addition, this interface can send or receive Hello packets. *Disable* – This parameter prevents this interface from communicating OSPF. In addition, this interface cannot send or receive Hello packets.
Multicast Forwarding	Not Supported.
Demand	Not Supported.
Transit Delay	Enter a value betweeen 0 and 3600 (the default value is 1). This value is the estimated number of seconds it takes to transmit a link-state update packet over this interface.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
Router Priority	Enter a value between 0 and 255. This number identifies the priority of the router associated with this logical port and is used to elect the designated and backup designated routers. The router with the highest priority is considered the designated router. A value of 0 indicates the router is not eligible to be the designated or backup designated router. If all routers have the same priority, the router ID is used to determine the designated router.
TOS 0 Metric	Enter a value between 1 and 65535. This value specifies the type of service cost. The lowest TOS 0 has the highest priority for routing.
Authentication Type	Specify the type of authentication that OSPF uses as a security measure to ensure that this logical port and router exchange information with correct neighbors. Options include: *None* – Specifies that no authentication is performed. *Simple Password* – Specifies a simple password authentication method that includes a password in all OSPF messages on an interface-by-interface basis. When a router receives a message on an interface that uses simple password authentication, the router checks the incoming OSPF message to see if the password is included in the message. If the password is correct, the message is processed normally. If the password is not part of the incoming message, the message is ignored and dropped. *MD5* – Specifies that an encryption method be used, which converts the authentication key to a number. The number is forwarded with the route rather than the actual key.
Authentication Key	Enter an authentication password in this field if you specified either *Simple* or *MD5* as the authentication type. This value is not required if you specified *None* as the authentication type.
Interval	
Re-Transmit	Enter a value between 0 and 3600 (the default value is 5 seconds). This value specifies the time to wait before resending a packet if no acknowledgment is received.
Hello	Enter a value between 1 and 65535 (the default value is 10 seconds). Specifies the number of seconds between router Hello messages. This parameter controls the frequency of router Hello messages on an interface.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
Router Dead	Enter a value greater than or equal to 0 (the default value for this field is 40 seconds). This value is a multiple of the Hello interval. For example, if the Hello interval is set to 10, the router dead interval should be configured at 20, 30, 40, etc. Specify this parameter if you have bad connections or if a link in the network is down. This parameter is the number of seconds a router waits to hear a Hello message from a neighbor before the router declares the neighbor "down." The value that you specify can affect OSPF operation. If the interval is too short, neighbors are considered down when they are reachable. If set for too long, routers that are really down are not considered down soon enough to properly reroute data.
Poll	Enter a value greater than or equal to 0 (the default value for this field is 120). Specifies the time, in seconds, between Hello packets sent to an inactive non-broadcast multi-access (NBMA) neighbor.
Operational Info (All values are read-only)	
Status	Displays the status of OSPF communication. *Options for Point-to-Point, Point-to-Multipoint, Broadcast, and Virtual link networks are:* *Up* – Indicates the network interface is operational. *Point-to-Point* – Indicates the interface is at the highest level of connection. In this state, the interface is operational and connects either to a physical point-to-point network or to a virtual link. Upon entering this state, the router attempts to form an adjacency with the neighboring router. Hello packets are sent to the neighbor at regular intervals based on the value that you specify for the *Hello Interval* parameter. See page 9-13 for details about the Hello Interval. *Init* – In this state, the neighbor sees a Hello packet. However, bidirectional communication has not been established with the neighbor. All neighbors in this state are listed in the Hello packets sent from the associated interface. *Down* – Indicates the interface is not usable. No protocol traffic will be sent or received on this interface.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
Status	*Options for an NBMA network are:* *Loopback* – In this state, the router's interface to the network is "looped back." The interface may be looped back in hardware and software. While in loopback, the interface is not available for regular traffic data traffic. *Waiting* – In this state, the router tries to determine the backup designated router's identity. To do this, the router monitors received Hello packets. The router cannot elect a backup designated router or designated router until it leaves the waiting state. This prevents any unnecessary changes to the backup designated router. *Designated Router* – In this state, the router is the designated router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate network link advertisements for the network node. The advertisement provides link information to all routers (including the designated router itself) attached to the network. *Backup Designated Router* – In this state, the router is the backup designated router on the attached network. When the present designated router fails, this router takes over. The router establishes adjacencies to all other routers attached to the network. *Other* – In this state, the router forms adjacencies to both the designated router and the backup designated router.
Designated Router	Displays the 32-bit IP address of the designated router for this network as seen by the advertising router. An IP address of 0.0.0.0 indicates that a designated router has not been specified for this network. If all routers have the same priority, the router ID is used to specify the designated router.
Backup Designated Rtr	Displays the 32-bit IP address of the backup designated router for this network as seen by the advertising router. An IP address of 0.0.0.0 indicates that a backup designated router has not been specified for this network.
Events	Displays the number of times this OSPF interface changed its state, or the number of times an error occurred.

4. When you are done setting parameters, choose OK.

Configuring OSPF Parameters at the Switch

This section describes how to configure the following OSPF switch parameters:

- IP parameters
- OSPF neighbors
- OSPF area aggregates
- OSPF virtual links

Configuring IP Parameters

Configure this parameter only if you configured the switch's OSPF interface in Area 1, which is used only for Ascend switches.

You also use Area 1 for routing updates between Ascend switches. This enables switches running switch software prior to 5.0 to operate with switches running 5.0 or later. If the interface is connected to a non-Ascend device, you cannot use Area 1.

To configure IP parameters:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set IP Parameters. The Set IP Parameters dialog box appears (Figure 9-5).

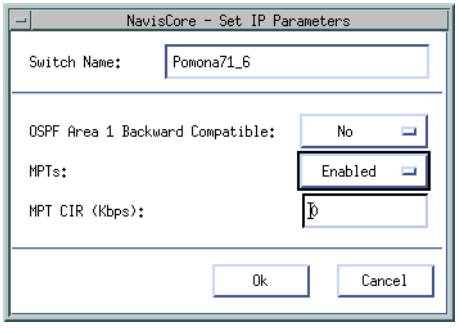

Figure 9-5. Set IP Parameters Dialog Box

3. Complete the fields described in Table 9-3.

Table 9-3. Set IP Parameters Field Descriptions

Field	Action/Description
Switch Name	Displays the name of the switch.
OSPF Area 1 Backward Compatible	Select either Yes or No. If you select Yes, the switch: – can communicate with other Ascend switches running pre-5.0 switch software – can communicate with other Ascend switches running 5.0 switch software, which is set to Yes in this field – *cannot* communicate with other vendor routers If you select No, the switch: – cannot communicate with other Ascend switches running pre-5.0 switch software – can communicate with other Ascend switches running 5.0 switch software, which is set to No in this field – *can* communicate with other vendor routers
MPTs	Set this value to Enable or Disable. In order for the switch to process MPTs, **the MPT administrative value for the switch must be set to Enable.** The MPT administrative value that you specify determines the use of MPTs on the switch as follows: • If the MPT value is set to Enable and no IP interfaces have been defined, the switch **does not establish MPTs.** • If the MPT value is set to Enable and IP interfaces have been defined, the switch **does establish MPTs** as a means of forwarding IP traffic. • If the MPT value is set to Disable and IP interfaces have been defined, the switch **does not establish MPTs to forward IP traffic,** but instead uses a hop-by-hop transmission method. • If the MPT value is set to Disable and no IP interfaces have been defined, the switch **does not establish MPTs.** Note: The NMS does not allow you to set an MPT administrative value to Disable on a switch that has one or more MPT point-to-point connections. You must delete all MPT point-to-point connections before you set the administrative value to Disable. See Chapter 12, "Multipoint-to-Point Tunneling," for more information.

Table 9-3. Set IP Parameters Field Descriptions (Continued)

Field	Action/Description
MPT CIR (Kbps) (Multi-Point-to-Point Tunneling Committed Information Rate)	Enter the rate in Kbps at which the Multipoint-to-Point tunnel (MPT) transfers data, averaged over a minimum increment of time. In addition, this value reserves bandwidth for MPTs, which the switch originates. For more information on MPTs, see Chapter 12, "Multipoint-to-Point Tunneling" ***Note:*** *This value applies to all links in the MPT.*

4. When you are done setting parameters, choose OK.

Defining an OSPF Neighbor

You do not have to define OSPF neighbors if you assign OSPF to an interface. OSPF automatically discovers its neighbors through Hello packets. However, if you configure an NBMA network, you must define OSPF neighbors. See the description for NBMA networks on page 9-12.

To define an OSPF neighbor:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Neighbors. The Set All OSPF Neighbors dialog box appears (Figure 9-6).

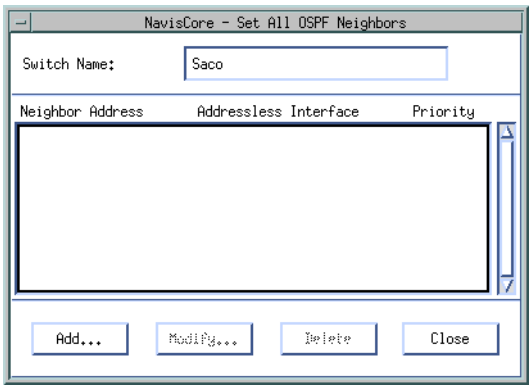

Figure 9-6. Set All OSPF Neighbors Dialog Box

The Set All OSPF Neighbors dialog box displays the following buttons:

Button	Function
Add	Enables you to add an OSPF neighbor.
Modify	Enables you to modify an OSPF neighbor.
Delete	Enables you to delete an OSPF neighbor.

3. Choose the Add button. The Add OSPF Neighbor dialog box appears (Figure 9-7).

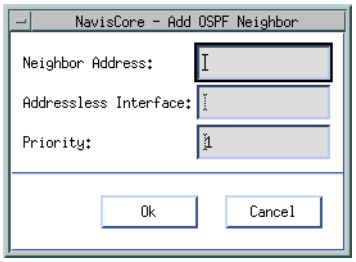

Figure 9-7. Add OSPF Neighbor Dialog Box

4. Complete the fields described in Table 9-4.

Table 9-4. Add OSPF Neighbor Fields

Field	Action/Description
Neighbor Address	Enter the IP address this neighbor uses in its IP source address. On address links, the address is not 0.0.0.0 but the address of the neighbor's interface.
Addressless Interface	Enter the addressless interface. If the interface has an IP address, the value is 0.0.0.0. If the interface is addressless, the value is the logical port number or interface number.
Priority	Enter a value between 0 and 255. The neighbor with the highest priority is the designated router. This field only applies to NBMA and broadcast networks. The value zero signifies the neighbor cannot be the designated router on this network.

5. When you are done setting parameters, choose OK.
6. At the Set All OSPF Neighbors dialog box, choose Close.

Defining an OSPF Area Aggregate

To define an OSPF area aggregate:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Area Aggregates. The Set All OSPF Area Aggregates dialog box appears (Figure 9-8).

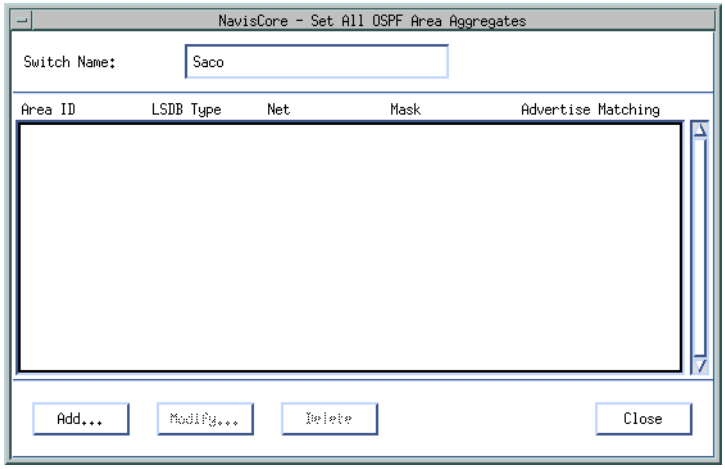

Figure 9-8. Set All OSPF Area Aggregates Dialog Box

The Set All OSPF Area Aggregates dialog box displays the following buttons:

Button	Function
Add	Enables you to add an OSPF area aggregate.
Modify	Enables you to modify an OSPF area aggregate.
Delete	Enables you to delete an OSPF area aggregate.

3. Choose the Add button.The Add OSPF Area Aggregate dialog box appears (Figure 9-9).

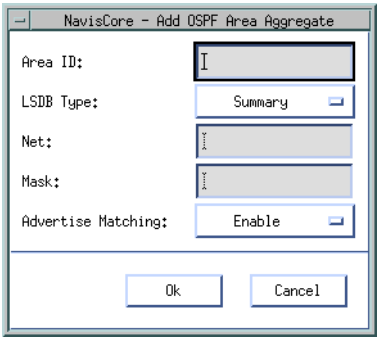

Figure 9-9. Add OSPF Area Aggregate Dialog Box

4. Complete the fields described in Table 9-5.

Table 9-5. Add OSPF Area Aggregate Fields

Field	Action/Description
Area ID	Enter the area ID (x.x.x.x) in which you want to locate the node. Area 0.0.0.0 is the network backbone. Areas are collections of networks, hosts, and routers. The area ID identifies the area. ***Note**: Area 1 is reserved for Ascend switches. If you configure the OSPF interface in Area 1, see "Configuring OSPF Parameters at the Switch" on page 9-16.*
LSDB Type	Specify the link state database type to which this address aggregate applies. Options include: *Summary* – Area border routers generate summary link advertisements, which describe inter-area routes (routes between areas) to networks. *NSSA External* – Not So Stubby Area external (NSSA) link advertisements allow an AS border router within a stub area and the routers within that area to learn about the external networks accessible through the AS border router in the area.
Net	Enter the IP address of the net or subnet, indicated by the range.
Mask	Enter the subnet mask that pertains to the net or subnet.
Advertise Matching	Select one of the following options: *Enable* – If you enable this parameter, you "leak" the net/mask you specified for the given area. *Disable* – If you disable this parameter, you hide the net/mask you specified for the given area.

5. When you are done setting parameters, choose OK.
6. At the Set All OSPF Area Aggregates dialog box, choose Close.

Defining an OSPF Virtual Link

To define an OSPF virtual link:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Virtual Links. The Set All OSPF Virtual Links dialog box appears (Figure 9-10).

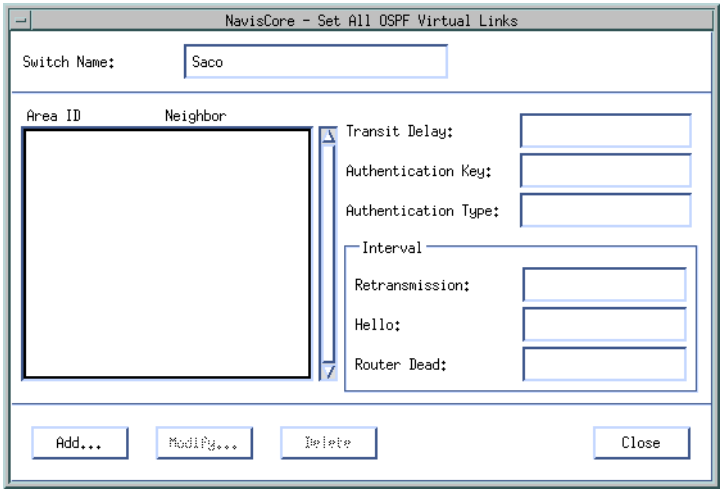

Figure 9-10. Set All OSPF Virtual Links Dialog Box

The Set All OSPF Virtual Links dialog box displays the following buttons:

Button	Description
Add	Enables you to add an OSPF virtual link.
Modify	Enables you to modify an OSPF virtual link.
Delete	Enables you to delete an OSPF virtual link.

3. Choose Add. The Add OSPF Virtual Link dialog box appears (Figure 9-11).

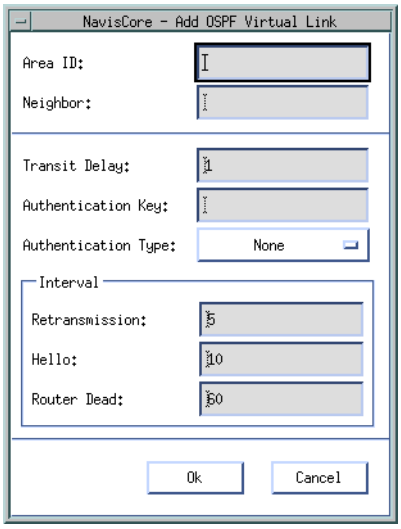

Figure 9-11. Add OSPF Virtual Link Dialog Box

4. Complete the fields described in Table 9-6.

Table 9-6. OSPF Virtual Link Fields

Field	Action/Description
Area ID	Enter the area ID (x.x.x.x) in which you want to locate the neighbor. Area 0.0.0.0 is the network backbone area. Areas are collections of networks, hosts, and routers. The area ID identifies the area. ***Note**: Area 1 is reserved for Ascend switches. If you configure the OSPF interface in Area 1, see "Configuring OSPF Parameters at the Switch" on page 9-16.*
Neighbor Address	Enter the IP address this neighbor uses in its IP source address. On addressless links, the address is not 0.0.0.0 but the address of the neighbor's interface.
Transit Delay	Enter a value between 0 and 3600 (the default value is 1). This field specifies the estimated number of seconds it takes to transmit a link-state update packet over this interface.
Authentication Key	Enter an authentication password in this field if you specified either *Simple* or *MD5* as the authentication type. This value is not required if you specified *None* as the authentication type.

Table 9-6. OSPF Virtual Link Fields (Continued)

Field	Action/Description
Authentication Type	Specify the type of authentication that OSPF uses as a security measure to ensure that this logical port and router exchange information with correct neighbors. Options include: *None* – Specifies that no authentication is performed. *Simple Password* – Specifies a simple password authentication method that includes a password in all OSPF messages on an interface-by-interface basis. When a router receives a message on an interface that uses simple password authentication, the router checks the incoming OSPF message to see if the password is included in the message. If the password is correct, the message is processed normally. If the password is not part of the incoming message, the message is ignored and dropped. *MD5* – Specifies that an encryption method be used, which converts the authentication key to a number. The number is forwarded with the route rather than the actual key.
Interval	
Retransmission	Enter a value between 0 and 3600 (the default value is 5 seconds) This field specifies the time to wait before resending a packet if no acknowledgment is received.
Hello	Enter a value between 1 and 65535 (the default value is 10 seconds). This field specifies the number of seconds between router Hello messages and controls the frequency of router Hello messages on an interface.
Router Dead	Enter a value greater than or equal to 0 (the default value for this field is 40 seconds). This value is a multiple of the Hello interval. For example, if the Hello interval is set to 10, the router dead interval should be configured at 20, 30, 40, etc. Specify this parameter if you have bad connections or if a link in the network is down. This parameter is the number of seconds a router waits to hear a Hello message from a neighbor before the router declares the neighbor “down.” The value that you specify can affect OSPF operation. If the interval is too short, neighbors are considered down when they are reachable. If set for too long, routers that are really down are not considered down soon enough to properly reroute data.

5. Choose OK.
6. At the Set All OSPF Virtual Links dialog box, choose Close.

Configuring an OSPF Route Map

Chapter 11, "Configuring Route Maps" provides detailed information about all types of route maps (including OSPF route maps) that you can configure, using IP services. See Chapter 11 before you begin any route map configuration

To configure an OSPF route map from the OSPF parameter menu:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Route Maps. The Set All OSPF Route Maps dialog box appears (Figure 9-12).

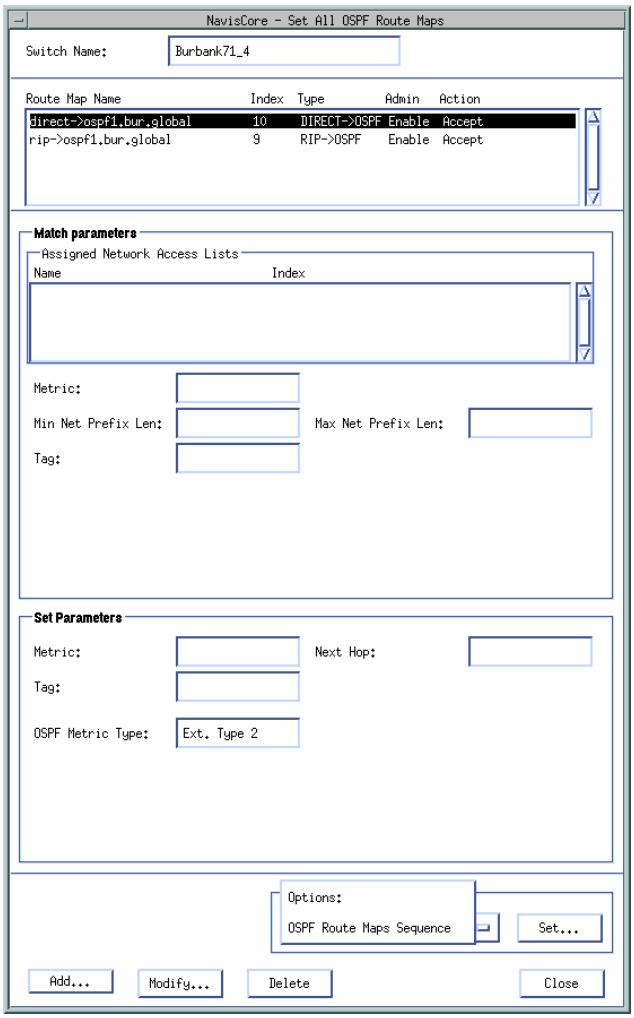

Use the OSPF Route Maps Sequence option and choose Set to display a dialog box that enables you to order the route map sequence.

Figure 9-12. Set All OSPF Route Maps

Configuring Multiple OSPF Areas

The configuration of multiple OSPF areas comprises the following configuration procedures:

- OSPF Area Configuration
- Virtual Link Configuration
- Address Aggregation

See "Steps for Configuring Multiple OSPF Areas" on page 9-27 for more information about each of these configuration procedures.

Area 0 is the OSPF backbone area. Areas do not have to be physically attached to the backbone. Instead, virtual links can be configured to logically attach an area to the backbone.

Every area border router must be connected to the backbone area. You use area 0 trunks or configured virtual links to connect each area border router to the backbone area.

Steps for Configuring Multiple OSPF Areas

The following sections outline each of the steps for OSPF Area configuration. This guide describes all configuration steps that reference each of the functions that you can access from the Ascend IP Parameters menu.

See the following guides for all steps that reference the functions that you access from the Ascend Parameters menu:

- *NavisCore Frame Relay Configuration Guide*
- *NavisCore ATM Configuration Guide*

Prerequisites

- The IP interface must be defined. See "Setting the IP Interface Address" on page 3-11 for details.
- Check the *NavisCore Frame Relay Configuration Guide* and the *NavisCore ATM Configuration Guide* for the necessary Frame Relay and ATM prerequisites.

Configuration Recommendations

Be aware of the following recommendations when configuring multiple OSPF areas:

- *Do not make areas too small.*
  - Area boundaries may be difficult to configure, and can cause sub-optimal routing for both IP and circuits.
  - If every switch is an area border switch, there will be no improvements to the route scaling.
- *Plan ahead when assigning areas.* Modification of the trunk Area ID is a service-affecting procedure that causes the trunk to bounce.
- *Do not unnecessarily aggregate IP addresses.* The process of modifying switch IP addresses is time-consuming and should not be done for the sole purpose of aggregation.

OSPF Area Configuration

1. Set the Area ID of all trunks. From the Administer menu, select Ascend Parameters ⇒ Set All Trunks. The NMS displays the NavisCore Set All Trunks dialog box. See the *Frame Relay Configuration Guide* or the *ATM Configuration Guide* for more information.
2. Set the Area ID of the IP logical ports. To do this:
   a. From the Administer menu, select Ascend IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.
   b. Choose IP Parameters. The Set IP Parameters dialog box appears.
   c. Choose IP Interface. The Set IP Interface Addresses dialog box appears.
   d. Choose Add OSPF. The Add OSPF Interface dialog box displays. See "Configuring OSPF at the Logical Port" on page 9-11 for more information about how to complete this dialog box.
3. Set the Area ID of the IP loopback addresses. To do this:
   a. From the Administer menu, select Ascend IP Parameters ⇒ Set All IP Loopback Addresses. The Set All IP Loopback Addresses dialog box appears. See "Setting IP Loopback Addresses" on page 8-16 for more information.
4. Set the Area ID of the Network Service Access Points (NSAPs). To do this:

a. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Node Prefixes. The Set All Node Prefixes dialog box appears. See the *NavisCore Frame Relay Configuration Guide* or the *NavisCore ATM Configuration Guide* for more information.

The Area ID of the switch IP address is set automatically to one of the following:

- Area 1, if it exists.
- If Area 1 does not exist, the Area ID is set to the ID of the switch with the lowest Area ID.

Virtual Link Configuration

Areas do not have to be physically attached to the backbone. Instead, you can configure virtual links to logically attach an area to the backbone. Before you can configure a virtual link you must know:

- The non-backbone transit area for the virtual link.
- The router IDs for the two endpoint switches. (The router ID of each switch endpoint is the same value as the switch ID).

In addition, you must ensure that both switches have IP addresses in the transit area. Add loopback addresses if necessary.

To add a virtual link:

1. From the Administer menu, select Ascend IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Virtual Links. The Set All OSPF Virtual Links dialog box displays (Figure 9-10 on page 9-23).
2. Choose Add. The Add OSPF Virtual Link dialog box appears (Figure 9-11 on page 9-24). Complete the fields that Table 9-6 on page 9-24 describes.

Address Aggregation

Aggregation is the process of advertising a single address prefix (rather than advertising multiple, more specific prefixes). Addresses can be aggregated at area borders. This practice further improves route scaling by reducing the size of the link-state database and the routing table.

You configure aggregates in the area border switch.

1. **To configure an aggregate for an NSAP,** from the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Node Prefixes. The Set All Node Prefixes dialog box appears.
2. **To configure an aggregate for an IP Address,** from the Administer menu, select Ascend IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Area Aggregates. The Set All OSPF Area Aggregates dialog box appears. See page 9-20 for details.

10

Configuring Static Routes

This chapter describes how to configure static routes.

About Static Routes

You configure static routes manually only if they are reachable. Static routes do not disappear from the IP routing table and will always be advertised. However, static routes do not respond to network topology changes. The only way a static route can change is if the network administrator changes them. In addition, static routes provide redundancy if a primary connection fails.

Configuring a Static Route

To configure a static route:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Static Routes. The Set All Static Routes dialog box appears (Figure 10-1).

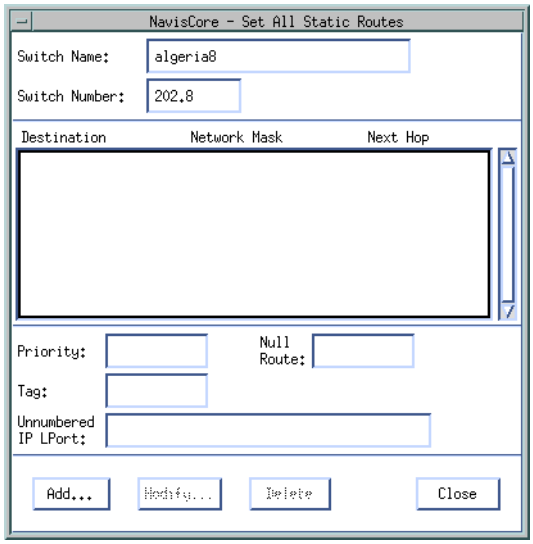

Figure 10-1. Set All Static Routes Dialog Box

The Set All Static Routes dialog box displays the following buttons.

Button	Function
Add	Enables you to add a static route.
Modify	Enables you to modify a static route.
Delete	Enables you to delete a static route.

3. Choose the Add button. The Set Static Route dialog box appears (Figure 10-2).

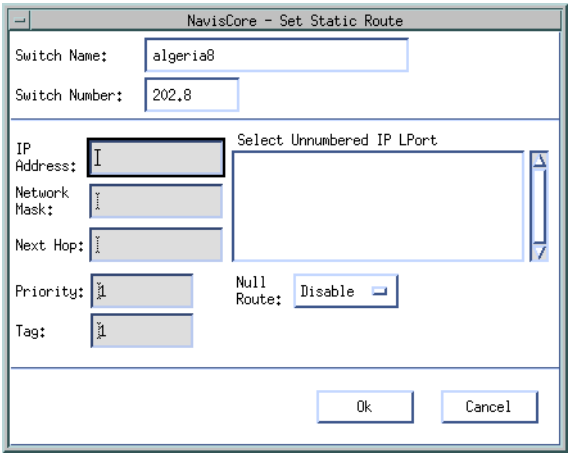

Figure 10-2. Set Static Route Dialog Box

4. Complete the values described in Table 10-1.

Table 10-1. Static Route Fields

Field	Description
Switch Name (read only)	Displays the name that identifies the switch.
Switch Number (read only)	Displays the switch number.
IP Address	Enter the IP address of the destination network.
Network Mask	Enter the network mask.
Next Hop	Enter the IP address of the next hop. The next hop field is disabled if you: – Selected an unnumbered IP logical port (see "Select Unnumbered IP LPort"), or – Enabled null route (see "Null Route")
Priority	Enter a value from 1 to 20 to specify the static route priority. The highest number is the preferred priority. The priority of the static route is in relation to other route protocols.
Tag	Enter the tag value, which you use to group multiple static route entries together.
Null Route	Select one of the following: *Enable* – If you enable this parameter, packets destined for this network will be discarded. In addition, the next hop is disabled. *Disable* – If you disable this parameter, packets destined for this network will be forwarded.
Select Unnumbered IP LPort	Select an unnumbered IP logical port to set up a static route to an IP interface that is not part of a subnet and does not have a specific address. Instead, the unnumbered IP logical port uses the router ID as its address.

5. Choose OK.
6. At the Static Route dialog box, choose Close.

11

Configuring Route Maps

This chapter describes the following configuration tasks:

- Adding a Network Filter
- Adding a Network Access List
- Adding a Route Map

About Route Maps

The purpose of a route map is to control and modify routing information and to define the parameters that your system uses to redistribute routes between routing domains. Route maps are used to alter route parameters that are then stored in the routing table, or sent via routing updates to other routers.

You can optionally define the following components for use in a route map:

- Network filters
- Network access lists

After you define the route map, you must assign it to a neighbor (in the case of BGP or RIP). If you are using multiple route maps for the same neighbor, you can specify the order that IP Navigator uses the specified maps.

The following sections define the concepts for using network filters, network access lists, and route maps. In addition, these sections describe how to use route maps to redistribute routes between routing domains.

About Network Filters

Network filters control the flow of route distribution. You can use a network filter to select routes that will be accepted or rejected by route maps. The specified filters must be used in a network access list and then applied to route maps.

When you create a network filter, you specify the following information:

- A network address
- A network mask value
- Coverage (inclusive or exact)

The network address and network mask value identify the route. The coverage specifies the type of access. *Inclusive* filters allow access to all networks that match the specified network address (including addresses that may be more specific such as a subnetwork address). *Exact* filters allow access only to the network that is specified in the network address.

A network filter is an optional component of a route map; however, if you want to use one or more network filters, you must include the filter in an access list and then include the access list in a route map. The route map must then be assigned to the appropriate neighbor or interface. A network filter by itself cannot be applied to a route map, neighbor, or interface.

About Access Lists

A network filter access list is an object that contains a set of unique network filters. Up to 300 network filters can be included in an access list.

You can create an empty network access list and later add defined network filters to the list. You use network access lists to logically group network filters.

A network access list is an optional component of a route map; however, if you want to use one or more network access lists, you must include the list in a route map and then assign the map to the appropriate neighbor or interface. A network access list by itself cannot be applied to a neighbor or interface

About Route Maps

Route maps enable you to specify the direction of route traffic based on the source of the traffic or a combination of both the traffic source and destination. You can enable or disable a route map as required by setting the Admin Status value for the route map.

When you create a route map, you specify two routing protocols: a From Protocol and a To Protocol. The route map specifies how routes are redistributed from one routing protocol to another. This is done between two different protocols as well as within the same protocol (for example, from BGP to BGP). The route maps are also used to selectively accept routes from a particular routing protocol into the router's main routing table.

In addition, you can optionally specify the following values as route map match or set parameters:

- Metric value
- Tag value
- Next hop address
- Autonomous System path values (BGP only)
- Community values (BGP only)
- OSPF route type

Route Map From and To Choices

Each time you define a route map, you must specify a From and a To choice to specify the two protocols used for route redistribution. The protocols that you specify govern the direction (import or export) as well as the set of affected routes.

The From choices include the following options:

- BGP
- OSPF
- RIP
- Static
- Direct
- Aggregate
- Any

The Any option enables you to select routes from the routing table regardless of the origin protocol. For example, you could select a specific route from the routing table and then advertise that route to BGP. The protocol used to transport the route to the routing table is not important.

The To choices that you can select vary depending upon the previously selected From choices. For example, the routing table option can only be selected if the From choice protocol is BGP or RIP.

The possible list of To choices include the following options:

- BGP
- OSPF
- RIP
- Routing Table

Determining if a Route Map is for Import or Export

The protocol that you specify for the To choice specifies whether a map is an import or export map as follows:

- Route maps that use a To choice of BGP, OSPF, or RIP are automatically created as export route maps.
- Route maps that use a To choice of Routing Table are created as import route maps.
- All route selections for a route map with the Routing Table as the To choice are performed before IP Navigator adds the routes to the routing table.

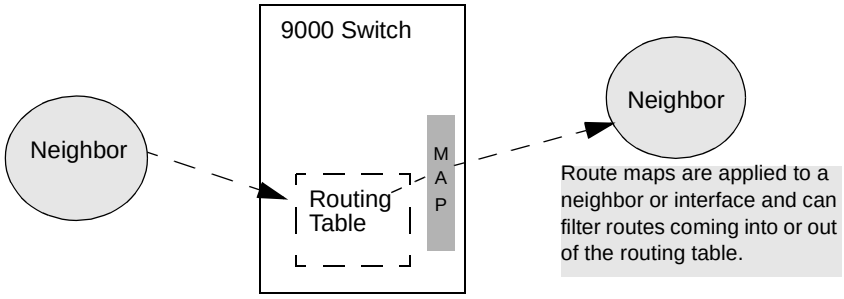

Figure 11-1. Using Route Maps to Filter Routes

Route Map Guidelines

Route maps are required if you want to accomplish any of the following tasks:

- Route filtering
- Route redistribution
- Altering route parameters such as metric, next hop, tag, and BGP path attributes.

See Figure 11-2 and the sections that follow for a description of the guidelines for route map use.

When are Route Maps Not Used?

You cannot use a route map to specify a routing policy for the following pairs:

OSPF to Routing Table — IP Navigator always adds OSPF routes to the routing table. For this reason, you cannot use a route map to specify the acceptance or rejection of specific routes between OSPF and the routing table.

OSPF to OSPF — IP Navigator always advertises OSPF routes to the OSPF routing domain. Link state protocols assume that all routes share the same information. For this reason, you cannot use a route map to specify the acceptance or rejection of specific routes being sent to an OSPF neighbor.

NMS Paths to OSPF — IP Navigator always advertises any NMS paths configured as Autonomous System External-Link State Advertisements (ASE-LSAs). (An NMS path is a static route that uses the Network Management Station as its destination.) For this reason, you cannot use a route map to specify the acceptance or rejection of specific NMS paths to the OSPF protocol.

Figure 11-2 illustrates the logical flow of routing information through the switch and where route maps can optionally be applied in this flow.

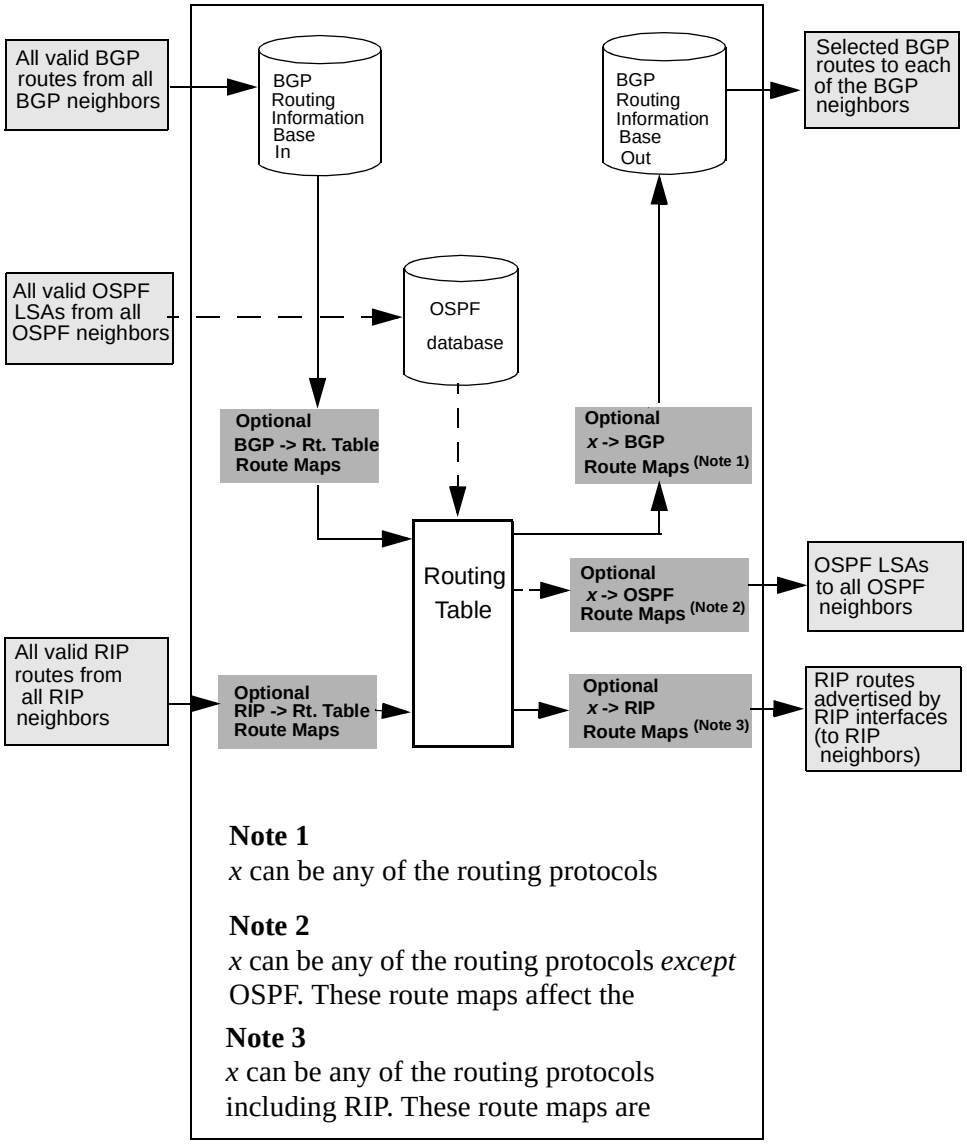

Figure 11-2. Flow of Routing Information Through the Switch

What Happens if You Do Not Use a Route Map?

If you do not use a route map for route filtering or route redistribution, the following import and export operations occur by default:

- Routes from all protocols, except for EBGP, are imported into the routing table by default.
- EBGP routes are not imported into the routing table by default for security reasons. You must specify a route map and optionally specify an access list containing any EBGP routes that you may want to import into the routing table.
- All RIP routes are exported to any RIP interface addresses that are configured for the IP interface.

Protocol Pairs That Do Not Require Route Maps

Route maps are not required for each of the following protocol pairs:

- IBGP Peer ⇒ Routing Table
- BGP ⇒ BGP
- RIP ⇒ Routing Table
- RIP ⇒ RIP

Protocol Pairs That Require Route Maps

Route maps are also required for each of the following protocol pairs:

Protocol	Description
Static ⇒ OSPF Direct ⇒ OSPF BGP ⇒ OSPF RIP ⇒ OSPF	Route maps are required in order to advertise any Static, Direct, BGP, and RIP routes into the OSPF routing domain. By default, IP Navigator does not advertise Static, direct, BGP, and RIP routes into the OSPF routing domain.
Static ⇒ RIP Direct ⇒ RIP BGP ⇒ RIP OSPF ⇒ RIP	Route maps are required to advertise any Static, Direct, BGP, and OSPF routes into the RIP routing domain. By default, IP Navigator does not advertise Static, Direct, BGP, and OSPF routes into the RIP routing domain.
Static ⇒ BGP Direct ⇒ BGP OSPF ⇒ BGP RIP ⇒ BGP	Route maps are required to advertise any Static, Direct, BGP, RIP, and OSPF routes into the BGP routing domain. By default, IP Navigator does not advertise Static, Direct, BGP, RIP, and OSPF routes into the BGP routing domain.
BGP ⇒ Routing Table	Route maps are required to install any routes advertised by neighboring EBGP peers into the main routing table. By default, IP Navigator does not install EBGP routes into the main routing table. IBGP routes are installed into the routing table even if there are no route maps.

IP Navigator applies multiple route maps using first match logic. This means that, as each route map is applied, any matching route entries are accepted or rejected immediately. *Subsequent route maps cannot consider the route entries that were already accepted or rejected.* For this reason, you should arrange the sequence of multiple route maps so that the *more specific matches are first in the list.*

Steps For Configuring a Route Map

To configure a route map, use the following steps:

1. *(Optional)* Define the network filters depending on your system's needs. See "Adding a Network Filter" on page 11-11 for more information.
2. *(Optional)* Use the defined network filters to create the network access lists. See "Adding a Network Access List" on page 11-13 for more information.
3. Specify the routing policies that define the match parameters to be used to filter routes and the set parameters for all selected routes. See "Adding Route Maps" on page 11-16 for more information.
4. Assign the route map to a BGP neighbor or a RIP interface. You assign route maps to BGP interfaces on the Modify BGP Neighbor dialog box. See Chapter 8, "Configuring BGP Parameters" for more information about accessing the BGP functions. You assign route maps to RIP interfaces on the Modify RIP Interface dialog box. See Chapter 7, "Configuring RIP" for more information about accessing the RIP functions.
5. If you are using multiple route maps, specify the order in which IP Navigator should use the assigned route maps by using the arrow buttons on the Modify BGP Neighbor and Modify RIP Interface dialog box. Route maps filter routes on the interface in the order in which they are specified on these dialog boxes. Route maps should be ordered from *most specific* to least *specific*.

Route maps that have a To protocol of OSPF are global and for this reason do not need to be assigned to an OSPF interface. IP Navigator uses this type of route map as soon as you create the map.

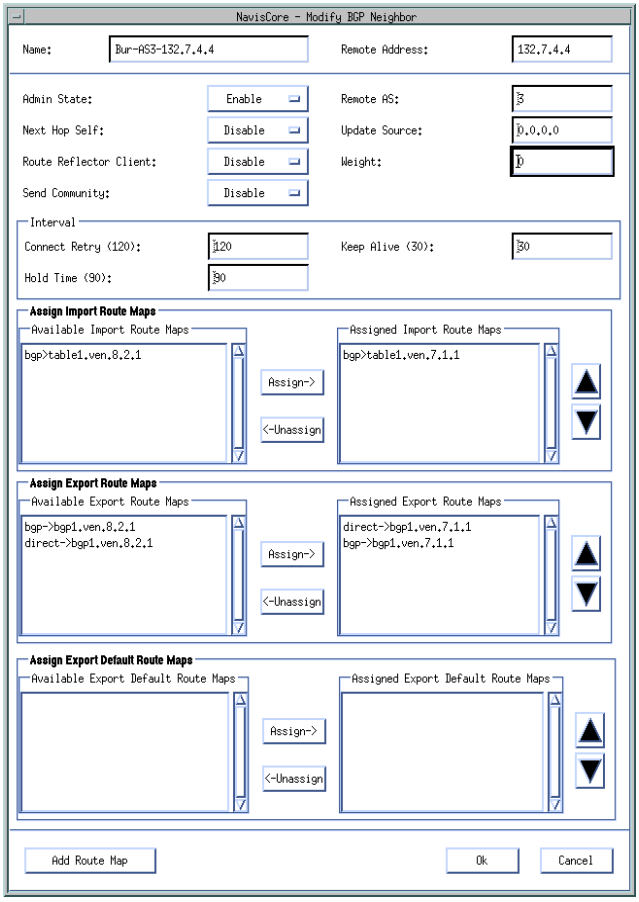

Use the arrow buttons to specify the sequence that IP Navigator uses to filter routing information.

Figure 11-3. Using the Arrow Buttons to Sequence Route Maps

Adding a Network Filter

To add a network filter:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Route Policies ⇒ Set All Network Filters. The Set All Network Filters dialog box displays.

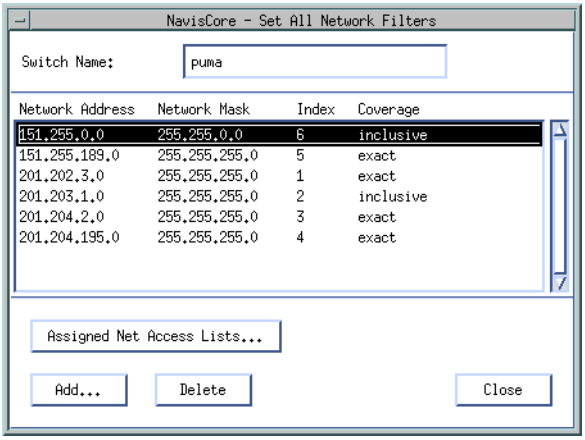

Figure 11-4. Set All Network Filters Dialog Box

Table 11-1 describes each of the Set All Network Filters buttons.

Table 11-1. Set All Network Filters Buttons

Button	Function
Add	Displays the Add Network Filter dialog box to enable you to add a network filter.
Delete	Displays the Delete Network Filter dialog box to enable you to delete a network filter.
Assigned Net Access Lists	Displays a list of any network access lists that use the selected filter.

3. Choose Add. The Add Network Filter dialog box displays (Figure 11-5).

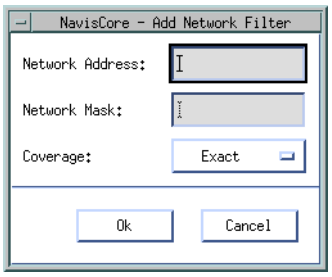

Figure 11-5. Add Network Filter Dialog Box

4. Specify the necessary network filter values listed in Table 11-2.

Table 11-2. Network Filter Fields

Field	Action/Description
Network Address	Specify the network address for this filter. For example, 0.0.0.0 specifies all network addresses.
Network Mask	Specify the network mask for this filter.
Coverage	Specify *inclusive* to allow all networks that match the specified network address (including addresses that may be more specific such as subnetwork addresses). Specify *exact* to allow only the network that is specified in the network address and the network mask.

Adding a Network Access List

A network access list enables you to logically group a set of network filters. To add a network access list:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Route Policies ⇒ Set All Network Access Lists. The Set All Network Access Lists dialog box displays.

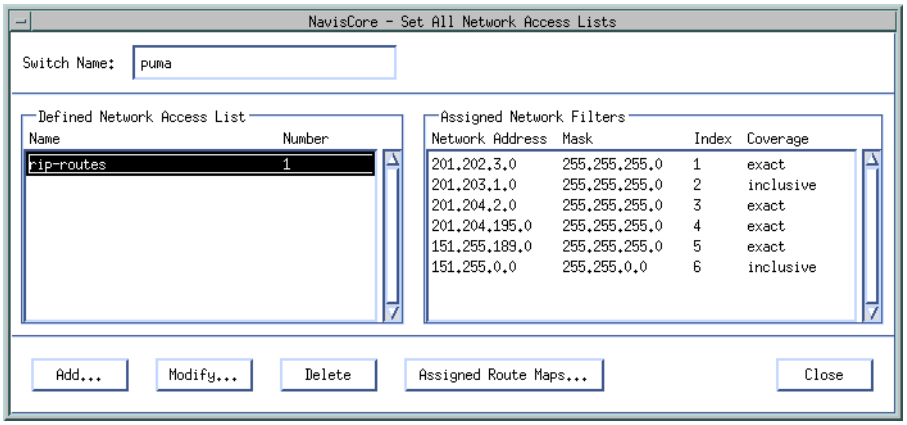

Figure 11-6. Set All Network Access Lists Dialog Box

Table 11-3 describes each of the Set All Network Access Lists buttons.

Table 11-3. Set All Network Access List Buttons

Button	Function
Add	Displays the Add Network Access List dialog box to enable you to add a network access list.
Modify	Displays the Modify Network Access List dialog box to enable you to modify a selected network access list.
Delete	Displays the Delete Network Access List dialog box to enable you to delete a selected network access list.
Assigned Route Maps	Displays a list of any route maps that use a selected network access list.

3. Choose Add. The Add Network Access List dialog box displays (Figure 11-7).

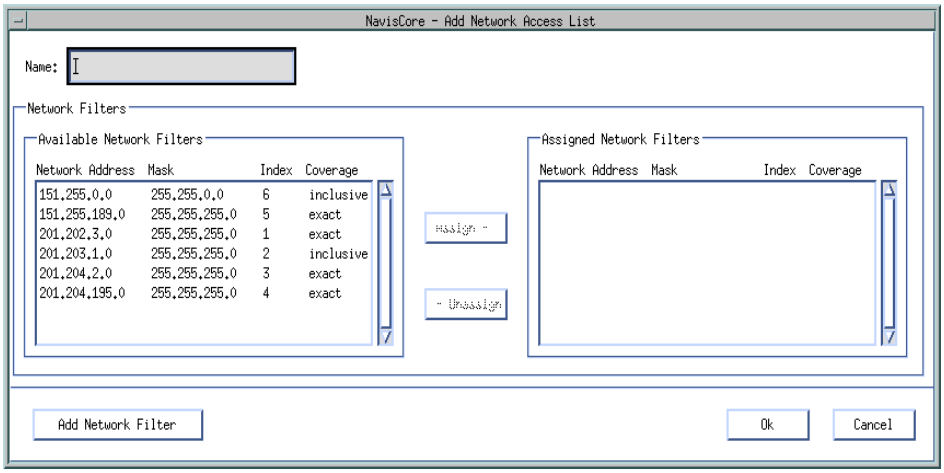

Figure 11-7. Add Network Access List Dialog Box

4. Specify a unique network access list name.
5. Use the Assign and Unassign buttons to specify the network filters that you want to include in the network access list. See Table 11-4 for a description of each of the fields on the Add Network Access List dialog box.
6. To add a filter to the list of Available Network Filters, choose Add Network Filter to display the Add Network Filter dialog box shown in Figure 11-5 on page 11-12. Any filters that you add are included in either the list of available network filters or the list of assigned network filters.
7. Choose OK after the Assigned Network Filters list includes all of the filters that you want to use in the network access list. One network access list can include up to 300 network filters.

Table 11-4. Network Access List Fields

Field	Action/Description
Name	Specify a unique network access list name.
Available Network Filters	A list of filters that are available for inclusion in the network access list.
Network Address	The network address for the filter.
Mask	The network mask for the filter.
Index	The index field is generated by NavisCore and is unique within the switch. This field is for internal system use only and cannot be modified.
Coverage	*Inclusive* allows all networks that match the specified network address (including addresses that may be more specific). *Exact* allows only the network that is specified in the network address.
Assigned Network Filters	A list of network filters that are currently included in the network access list. Up to 300 filters can be included in the access list.
Network Address	The network address for the filter.
Mask	The network mask for the filter.
Index	The index field is generated by NavisCore and is unique within the switch. This field is for internal system use only and cannot be modified.
Coverage	*Inclusive* allows all networks that match the specified network address (including addresses that may be more specific). *Exact* allows only the network that is specified in the network address.

Adding Route Maps

To add a route map:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All Route Policies ⇒ Set All Route Maps. The Set All Route Maps dialog box displays.

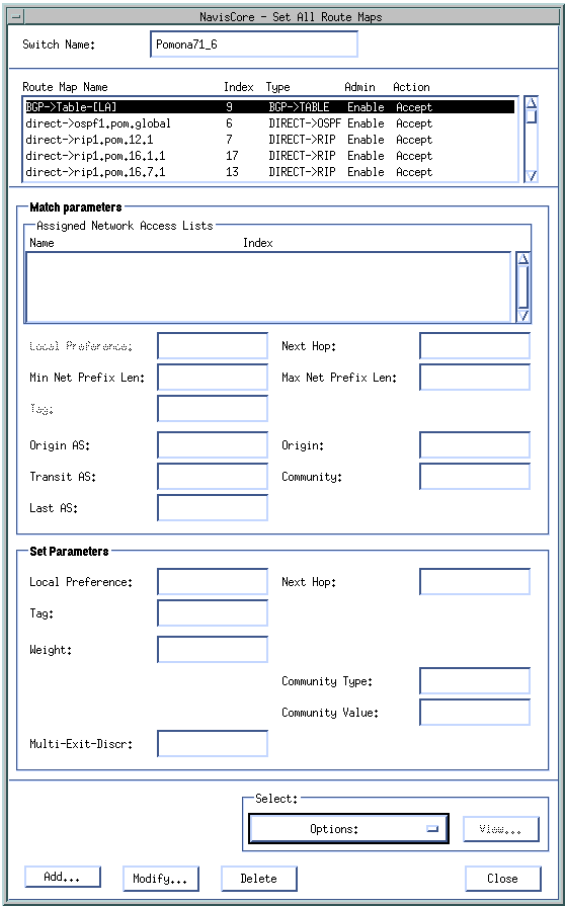

Figure 11-8. Set All Route Maps Dialog Box

Table 11-5 describes each of the Set All Route Maps buttons. Table 11-6 describes the fields at the top of the Set All Route Maps dialog box.

The Match Parameters and Set Parameters on the Set All Route Map dialog box vary depending on the type of route map that you are defining. See Table 11-9 for a reference to the section of this chapter that describes the Match and Set parameters for each route map type.

Table 11-5. Set All Route Maps Buttons

Button	Function
Add	Displays the Add Route Map dialog box to enable you to add a route policy.
Modify	Displays the Modify Route Map dialog box to enable you to modify a route policy.
Delete	Displays the Delete Route Map dialog box to enable you to delete a route policy.
Options	 Use the Select: Options button to select the following option. **Assigned BGP Neighbors** — Lists all BGP neighbors that use a selected route map. **Assigned RIP Interfaces** — Lists all RIP interfaces that use a selected route map. **BGP Neighbors** — Displays the Set All BGP Neighbors dialog box to enable you to assign a route map to a BGP neighbor. **OSPF Route Maps Sequence** — Displays the Change the Order of OSPF Route Maps dialog box to enable you to change the sequence of assigned route maps.

Table 11-6. Set All Route Maps Common Values

Field	Action/Description
Switch Name	Displays the name of the currently selected switch.
Route Map Name	A name that uniquely identifies the route map.
Index	The index field is generated by NavisCore and is unique within the switch. This field is for internal system use only and cannot be modified.
Type	Displays the From protocol and To protocol that identify the route distribution type. See Table 11-9 on page 11-21 for a list of route distribution types and a reference to the section of this chapter that describes how to redistribute routes between various routing protocols.
Admin	Specify Enable or Disable. *Enable* indicates that the route map is administratively enabled and can be used. *Disable* indicates that the route map is administratively disabled and cannot be used.
Action	Specify Accept, Deny, or Originate Default. *Accept* – indicates that all routes that match the specified Match parameters are accepted. *Deny* – indicates that all routes that match the specified Match parameters are denied. *Originate Default* – indicates that you can specify the match parameters that define where to send a default route heading. This option is used for the following types of route maps: BGP to BGP, ANY to BGP, or RIP to RIP.

3. Choose Add. The Add Route Map dialog box displays.

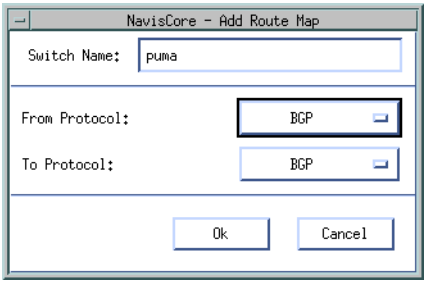

Figure 11-9. Add Route Map Dialog Box

4. Specify the values listed in Table 11-7.

Table 11-7. Route Map Descriptions

Field	Action/Description
Switch Name	Displays the name of the currently selected switch.
From Protocol	Specify one of the following values: BGP, OSPF, RIP, STATIC, Direct, Aggregate, or ANY.
To Protocol	Specify one of the following values: BGP, OSPF, RIP, or Routing Table. The routing table option can only be selected if the From protocol is either BGP or RIP.

If you configure a route map and specify ANY or DIRECT as the From Protocol, make sure that you also configure an access list that selects only those routes that you want to include as export routes.

5. Choose OK. The system displays a dialog box similar to the one shown in Figure 11-10.

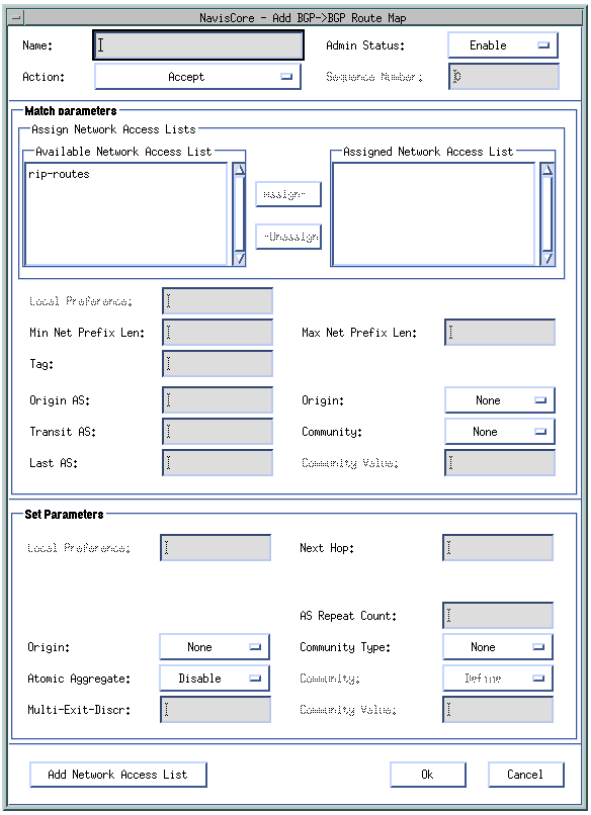

Figure 11-10. Second Add Route Map Dialog Box

6. Specify the Route Map Name, Admin Status, and Action values as described in Table 11-8.
7. Specify the necessary match and set parameters for this route map. If you need to add an access list for inclusion in the route map, choose Add Access Lists. Instructions for adding access lists start on page 11-13.

   The Match parameters and Set parameters on the Add Route Map dialog box vary depending on the type of route map that you are defining. See Table 11-9 for a reference to the section of this chapter that describes the Match and Set parameters for each route map type.

All of the Match and Set Parameter fields described in Table 11-10 through Table 11-25 are optional. It is possible to specify a routing policy that uses no match and no set values.

Table 11-8. Add Route Map Fields

Field	Action/Description
Route Map Name	Specify a unique name to identify the route map.
Admin Status	Specify Enable or Disable. *Enable* indicates that the route map is administratively enabled and can be used. *Disable* indicates that the route map is administratively disabled and cannot be used.
Action	Specify Accept, Deny, or Originate Default. *Accept* – indicates that all routes that match the specified Match parameters are accepted. *Deny* – indicates that all routes that match the specified Match parameters are denied. *Originate Default* – indicates that you can specify the match parameters that define where to send a default route heading. This option is used for the following types of route maps: BGP to BGP, ANY to BGP, or RIP to RIP.

Table 11-9. Match and Set Parameter Field Descriptions

Route Map Type	See...
BGP to BGP	Table 11-10 on page 11-22
BGP to OSPF	Table 11-11 on page 11-25
BGP to RIP	Table 11-12 on page 11-27
BGP to Routing Table	Table 11-13 on page 11-29
OSPF to BGP	Table 11-14 on page 11-32
OSPF to RIP	Table 11-15 on page 11-34
RIP to RIP	Table 11-16 on page 11-35
RIP to BGP	Table 11-17 on page 11-36
RIP to OSPF	Table 11-18 on page 11-38
RIP to Routing Table	Table 11-19 on page 11-39
Static to OSPF	Table 11-20 on page 11-40
Static to BGP	Table 11-21 on page 11-41
Static to RIP	Table 11-22 on page 11-43
Any or Direct to BGP	Table 11-23 on page 11-44
Any or Direct to OSPF	Table 11-24 on page 11-46
Any or Direct to RIP	Table 11-25 on page 11-47
Aggregate to BGP	Table 11-26 on page 11-48

If you configure a route map and specify ANY or DIRECT as the From protocol, make sure that you also configure an access list that selects *only those routes* that you want to include as export routes.

Table 11-10. BGP to BGP Match and Set Parameter Fields

Field	Description
Match Parameters	BGP routes can be distributed to BGP based on matches to the following parameters. *Any fields that you do not plan to use as a match parameter should be left blank.*
Assign Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	The route tag value. Tag values are used to further identify a route. Only routes matching the specified tag value are selected.
Origin AS	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. 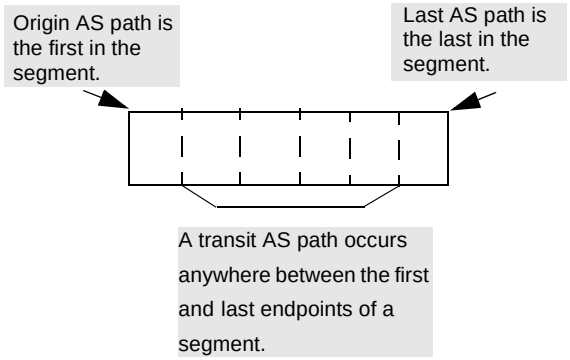
Transit AS	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.
Last AS	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.
Origin	Specify one of the following values to indicate the BGP origin code for use as a match parameter: *IGP, EGP, Incomplete, None.*

Table 11-10. BGP to BGP Match and Set Parameter Fields (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank. No default is used if the field is left blank.*
Origin	Specify one of the following values to indicate the BGP origin code for use as a match parameter: *IGP, EGP, Incomplete, None.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute should be set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values to identify the community type: *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification will occur.

Table 11-10. BGP to BGP Match and Set Parameter Fields (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* –Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-11. BGP to OSPF Match and Set Parameter Fields

Field	Description
Match Parameters	BGP routes can be distributed to OSPF based on matches to the following parameters. Only routes that match the specified parameters are selected for the Set operations. *Any fields that you do not plan to use as a match parameter should be left blank.*
Assign Network Access Lists	Use the *Assign* and *Unassign* options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value. Tag values are used to further identify a route. Only routes matching the specified tag value will be selected.
Origin AS	Specify a match parameter for the Autonomous System (AS) where the route originated.
Transit AS	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route.
Last AS	Specify a match parameter for the last Autonomous System (AS) in the route.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, or None.*
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-11. BGP to OSPF Match and Set Parameter Fields (Continued)

Field	Description
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the OSPF route metric to the specified metric value. If you leave this field blank, a default metric from the routing table is used.
Tag	Sets the OSPF route tag value to the specified value. If you leave this field blank, a default tag from the routing table is used.
OSPF Metric Type	Specify *External-type-1* or *External-type-2.* If you leave this field blank, External-type-2 is used as the default.
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default of 0 is used.

Table 11-12. BGP to RIP Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from BGP to RIP are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Origin AS	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. 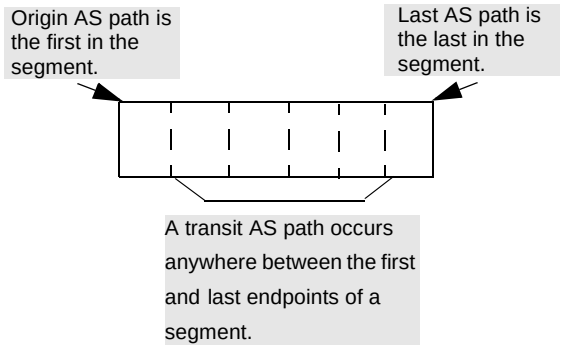
Transit AS	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.
Last AS	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete,* or *None.*

Table 11-12. BGP to RIP Match and Set Parameter Fields (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export*, *No Advertise*, or *Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the RIP route metric to the specified metric value. If you leave this field blank, a default metric from the routing table is used.
Tag	Sets the route tag field for the route. If you leave this field blank, a default tag from the routing table is used.
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default of 0 is used.

Table 11-13. BGP to Routing Table Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from BGP to the routing table is based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the Assign and Unassign options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Origin AS	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. 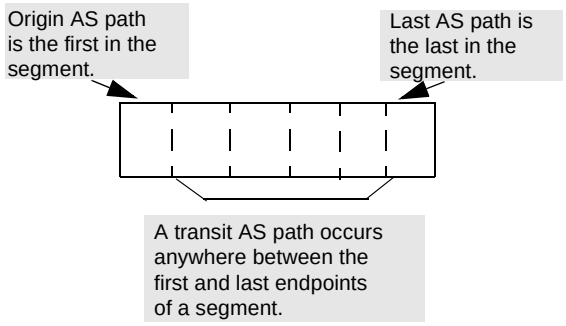
Transit AS	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.
Last AS	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.
Next Hop	Specify the IP address that specifies the next hop to reach a network. Only routes that match this next hop value are selected.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete,* or *None.*

Table 11-13. BGP to Routing Table Match and Set Parameter Fields (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export*, *No Advertise*, or *Local AS*. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Tag	Sets the route tag field for the route.
Weight	A weight value that is assigned to a route. This value is used only for routes from EBGP peers. The weight value is not used for routes from IBGP peers.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that specifies the next hop to reach a network.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. Additive – Adds a community to an existing community. *None* – No community modification occurs.

Table 11-13. BGP to Routing Table Match and Set Parameter Fields (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-14. OSPF to BGP Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from an OSPF domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The OSPF cost. If you leave this field blank, a default value from the routing table is used.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
OSPF Route Type	Specify one of the following OSPF Metric Type values: *Intra, Internal, External-1, External-2,* or *None.*
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that specifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.

Table 11-14. OSPF to BGP Match and Set Parameter Fields (Continued)

Field	Description
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. Additive – Adds a community to an existing community. *None* – No community modification occurs.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-15. OSPF to RIP Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from OSPF to RIP are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The OSPF cost. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
OSPF Route Type	Specify one of the following OSPF Metric Type values: *Intra, Internal, External-1, External-2,* or *None*.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The route tag field for the route that you want to set. If you leave this field blank, a default tag from the routing table is used.
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-16. RIP to RIP Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from RIP or RIP version 2 to RIP or RIP version 2 are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The route tag field for the route that you want to set. If you leave this field blank, a default tag from the routing table is used.
Next Hop	An IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-17. RIP to BGP Match and Set Parameter Fields

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into a BGP domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the origin of the route: *IGP, EGP, Incomplete,* or *None.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that specifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.

Table 11-17. RIP to BGP Match and Set Parameter Fields (Continued)

Field	Description
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: *No Export, No Advertise,* or *Local AS.* *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-18. RIP to OSPF Match and Set Parameter Fields

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into an OSPF domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to OSPF. If you leave this field blank, a default tag from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. If you leave this field blank, a value of External Type 2 is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-19. RIP to Routing Table

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into an OSPF domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Next Hop	Specify the IP address that specifies the next hop to reach a network. Only routes that match this next hop value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The tag to be set in the redistributed routes. If you leave this field blank, a default tag from the routing table is used.

Table 11-20. Static to OSPF Match and Set Parameter Fields

Field	Description
Match Parameters	Static routes can be distributed to OSPF based on matches to the following lists. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The OSPF cost. If no OSPF metric is specified, a default metric value from the routing table is used.
Tag	The tag to be set in the redistributed routes to OSPF. If none is specified, then a default tag value from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2.*
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-21. Static to BGP Match and Set Parameters

Field	Description
Match Parameters	Static routes can be distributed to BGP based on matches to the following parameters. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the Assign and Unassign options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes to the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the origin of the route: *IGP, EGP, Incomplete,* or *None.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that specifies the next hop to reach a network. The next hop value is set to this value on all selected routes.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs.

Table 11-21. Static to BGP Match and Set Parameters (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export*, *No Advertise*, or *Local AS*. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-22. Static to RIP Match and Set Parameter Fields

Field	Description
Match Parameters	Static routes can be distributed to RIP based on matches to the following lists. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the metric value on all selected routes to the specified metric value. If you leave this field blank, a default metric value from the routing table is used.
Tag	Sets the tag value on all selected routes to the specified tag value. If you leave this field blank, a default tag value from the routing table is used.
Next Hop	An IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-23. Any or Direct to BGP

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The OSPF cost. Only routes matching this value are selected. This parameter is not used for Direct to BGP route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the OSPF route tag value to be used as the match parameter. Only routes matching this value are selected. This parameter is not used for Direct to BGP route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as a set parameter should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes to the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify Enable or Disable to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	An IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.

Table 11-23. Any or Direct to BGP (Continued)

Field	Description
Community Type	Specify one of the following values. *Replacement* –Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs. All selected routes are set to the value that you specify.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise,* or *Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-24. Any or Direct to OSPF Parameters

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The OSPF cost. Only routes matching this value are selected. This parameter is not used for Direct to OSPF route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected. This parameter is not used for Direct to OSPF route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to OSPF. If no value is specified, a tag value from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. The OSPF Metric Type on selected routes is set to the specified value. A default value of *External Type 2* is used if no value is specified.
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-25. Any or Direct to RIP Parameters

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into RIP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The metric value that is used as a match parameter. Only routes matching this value are selected. This parameter is not used for Direct to RIP route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected. This parameter is not used for Direct to RIP route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If no RIP metric is specified, the value in the routing table is used.
Tag	The tag to be set in the redistributed routes to RIP. If no value is specified, a tag value from the routing table is used.
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-26. Aggregate to BGP

Field	Action/Description
Match Parameters	There are no match parameters for an Aggregate to BGP route map.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as a set parameter should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes to the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify Enable or Disable to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that specifies the next hop to reach a network. Only routes that match this next hop value are selected.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number is assigned to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs. All selected routes are set to the value that you specify.

Table 11-26. Aggregate to BGP (Continued)

Field	Action/Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: *No Export, No Advertise,* or *Local AS.* *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

12

Multipoint-to-Point Tunneling

This chapter provides an overview of Multipoint-to-Point Tunnels (MPTs) and MPT Point-to-Point connections. The following MPT configuration tasks are also described:

- Configuring an ATM OPTimum cell trunk for MPT traffic
- Configuring an MPT Point-to-Point connection
- Defining an MPT Point-to-Point connection path
- Displaying the operational status for an MPT Point-to-Point connection path

About MPTs

IP Navigator uses MPTs as a means of forwarding IP traffic over switched paths (no intermediate IP lookups) through the Ascend network. An MPT allows multiple nodes to share the same circuit for transmission to a single destination. An MPT can be thought of as the inverse of the point-to-multipoint virtual circuit used to allow the sending of packets from one source to multiple destinations. Connections of this type are commonly used in multicast applications such as video distribution. Because all nodes in a network share the same circuit to transmit to a specific destination node, the number of circuits (MPTs) needed to fully interconnect all of the nodes is approximately equal to the number of nodes.

Every switch maintains one MPT circuit network. This MPT network is *rooted* at the switch. For this purpose, the switch maintains a root, which

- keeps track of MPT nodes
- adds and deletes nodes
- keeps nodes alive

A root is a standard circuit endpoint that is created at initialization time on every CP card in the 9000 and every SP card in the 500. All other nodes on this MPT circuit network are considered leaves. As noted before, traffic flow occurs from the leaves to the root.

MPT Administrative Value

In order for the switch to process MPTs, **the MPT administrative value for the switch must be enabled.** You set the MPT administrative value on the Set IP Parameters dialog box (Figure 9-5 on page 9-16). The MPT administrative value that you specify determines the use of MPTs on the switch as follows:

- If the MPT value is set to Enable and no IP interfaces have been defined, the switch **does not establish MPTs.**
- If the MPT value is set to Enable and IP interfaces have been defined, the switch **does establish MPTs** as a means of forwarding IP traffic.
- If the MPT value is set to Disable and IP interfaces have been defined, the switch **does not establish MPTs to forward IP traffic,** but instead uses a hop-by-hop forwarding method.
- If the MPT value is set to Disable and no IP interfaces have been defined, the switch **does not establish MPTs.**

The NMS does not allow you to set an MPT administrative value to Disable on a switch that has one or more MPT point-to-point connections. You must delete all MPT point-to-point connections before you set the administrative value to Disable. See "MPT Point-to-Point Connections" on page 12-7 for more information.

Switch Domains

There are two types of switch domains:

Cell Domain — Paths that traverse direct ATM trunks.

Frame Domain — Paths that traverse pure direct frame trunks.

A switch can belong to multiple domains, however, the domains must be adjacent. Switches that belong to multiple domains must reside at the border of these domains. In addition, these switches must perform additional protocol layer processing to determine routes across the different domains.The root maintains connections to each domain the switch belongs to.

OSPF determines how MPTs connect two switches in different domains. The following factors apply when determining MPTs:

- A switch that only belongs to one domain cannot add a switch from a different domain to its MPT. **MPTs are only established between switches in the same domain.**
- If the shortest path between two switches in the same domain traverses a different domain, the switches cannot add each other to their MPTs.

MPTs use cell and frame domains to circumvent addressing limitations in switching ATM cells. **Whenever you cross boundaries between cell and frame domains, an IP lookup is required.**

OSPF Areas

MPTs are limited to a single OSPF area. Within an area, IP traffic can be switched, however, at area borders IP traffic must be routed. When traffic is routed at the area boundaries, the IP header must be examined and a routing table lookup is made. See Chapter 9, “Configuring OSPF Parameters” for a detailed description of OSPF areas.

Processing MPTs

MPTs on the CBX 500 are used to forward IP data over virtual paths (VPs) from one switch to another. MPTs are initiated on the SP of a CBX 500 in the same way that they are initiated at the CP on a B-STDX 8000/9000.

However, each leaf that is added to the MPT occurs:

- in the CP on a B-STDX 8000/9000.
- in the SP on a CBX 500. Each 9000 node and the SP and each forwarding engine (FE) on a CBX 500 is added as a leaf of the MPT. FEs reassemble cells and perform IP lookups. There are two FEs in each of the following CBX 500 cards:
  - 6-port DS3 Frame card
  - 4-port Ethernet

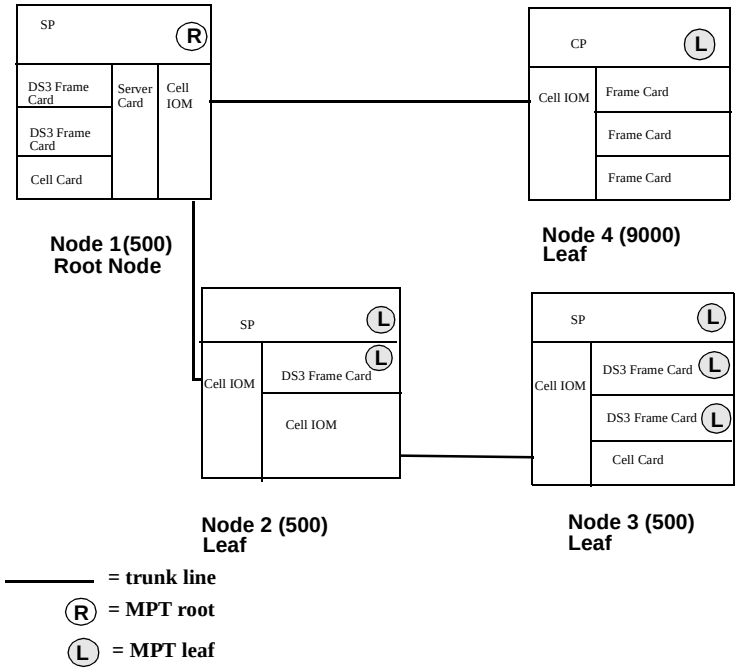

Figure 12-1. MPT Leaf Occurrences in the CBX 500 and B-STDX 8000/9000

MPTs Over OPTimum Trunks

An OPTimum trunk creates a switch-to-switch Ascend trunk through a Public Data Network (PDN). The Ascend OPTimum trunk feature allows private enterprise networks to purchase lower-cost, public-carrier services and configure OPTimum trunks between two Ascend switches. For more information about configuring an OPTimum trunk see the following guide:

- *NavisCore ATM Configuration Guide*

IP Navigator automatically assigns a Permanent Virtual Path (PVP) to each MPT crossing an OPTimum trunk. These point-to-point PVPs carry MPT traffic in both directions and, therefore reduce the number of paths required to interconnect switches in two given clusters.

Before IP Navigator can assign PVPs, you must specify a range of VPIs for each logical port endpoint of the OPTimum trunk. To do this, you must specify the VPI range on the OPT Trunk VPI Range attributes dialog box.

Configuring an ATM OPTimum Cell Trunk for MPT Traffic

Use the following steps to configure an ATM OPTimum Cell Trunk for MPT traffic:

1. Access the Add Logical Port dialog box. For complete details about how to access the Add Logical Port dialog box and information about provisioning OPTimum cell logical ports and trunks, see the *NavisCore ATM Configuration Guide.*
2. Access the Opt Trunk VPI Range attributes on the Set Attributes menu. NavisCore displays the dialog box shown in Figure 12-2.

   For complete details about how to access the Set Attributes menu, see the *NavisCore ATM Configuration Guide.*

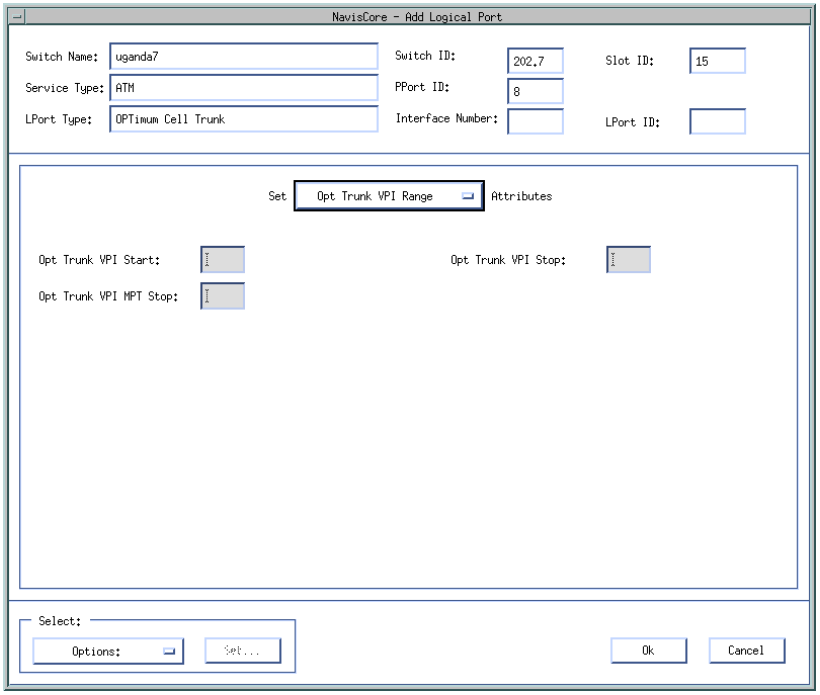

Figure 12-2. Add Logical Port – Opt Trunk VPI Range Dialog Box

3. Specify the Opt Trunk VPI Start and Stop Values. The VPI start and stop values specify the range of VPIs that can be created over this OPTimum trunk. The maximum allowable range is a value from 1 through 255 for a UNI logical port and 1 through 4095 for an NNI logical port. Since you can specify more than one OPTimum trunk on the same physical link, make sure that the VPI value on each trunk does not exceed these limits.

4. Specify the OPT Trunk VPI MPT Stop value. To use an ATM OPTimum cell trunk for MPT traffic, enter a VPI MPT stop value that specifies the part of the VPI range that is dedicated to MPT traffic. For example:

VPI start = 2
VPI MPT stop = 10
VPI Stop = 15

VPIs 3 through 10 are dedicated to MPT traffic only; VPIs 2, and 11 through 15 are used for other VPs.

The default value of 0 prevents MPT traffic from using the OPTimum cell trunk.

The range of VPIs configured for MPT traffic must be identical on both sides of the trunk.

MPT Point-to-Point Connections

Switches in an IP Navigator network automatically establish MPT circuits upon startup. MPT Point-to-Point connections are user-defined circuits for IP traffic between exactly two switches. This feature enables you to provide more bandwidth between two nodes. All traffic will use the user-defined point-to-point connection rather than the automatic connection.

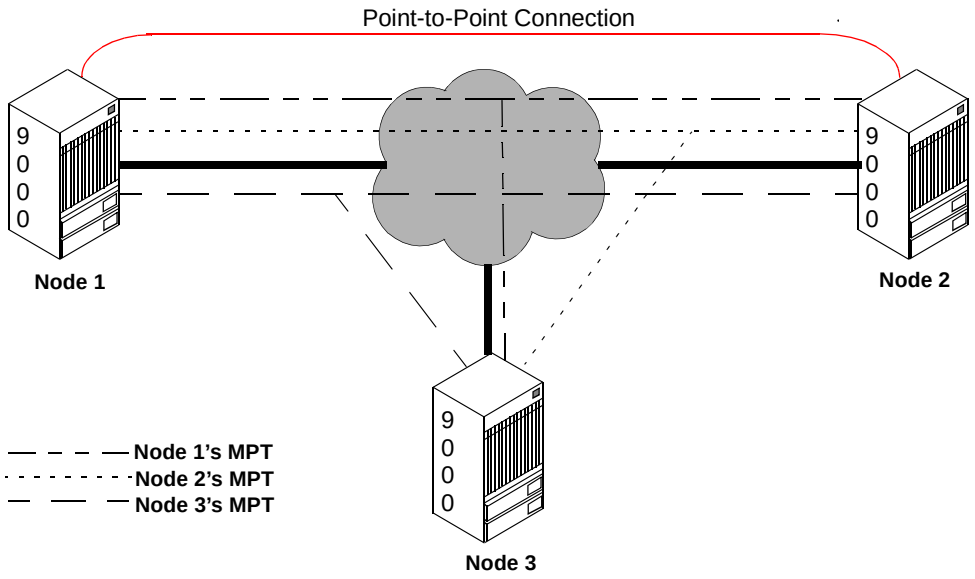

Figure 12-3. MPT Point-to-Point Connections

You can define an MPT Point-to-Point connection between:

- Two B-STDX 8000/9000 switches
- A CBX 500 and a B-STDX 8000/9000 switch, if the CBX 500 switch has a DS3 Frame card or if the B-STDX 8000/9000 switch has a cell card. (Cell cards include any of the following cards: the ATM T3 UNI card, the ATM-CS DS3/E3 card, and the ATM IWU/OC3 card.)
- Two CBX 500 switches.

Configuring an MPT Point-to-Point Connection

To configure an MPT Point-to-Point Connection from the NavisCore menu:

1. Select Ascend IP Parameters ⇒ Set MPT Point-to-Point from the Administer menu. The Set All MPT Point-to-Point Connections dialog box appears (Figure 12-4). Table 12-1 on page 12-8 describes each of the Set All MPT Point-to-Point Connections fields.

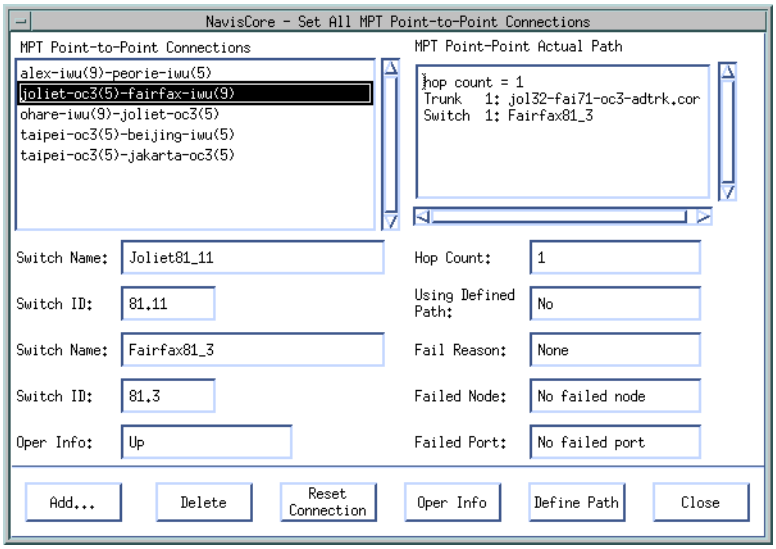

Figure 12-4. Set All MPT Point-to-Point Connections Dialog Box

Table 12-1. Set All MPT Point-to-Point Connections Buttons

Button	Description
Add	Displays the Select MPT Point-to-Point Connection Endpoints dialog box (Figure 12-5) to enable you to specify the endpoints for an MPT point-to-point connection.
Delete	Deletes a selected MPT point-to-point connection.
Reset Connection	Enables you to reestablish the connection signalling in the event of a failure.
Oper Info	Displays operational information about a selected MPT point-to-point connection.
Define Path	Displays the Set MPT Point-to-Point Defined Path dialog box to enable you to define the path of a selected MPT point-to-point connection. See “MPT Point-to-Point Connections” on page 12-7.

Table 12-2. Set All MPT Point-to-Point Connections Field Descriptions

Field	Action/Description
MPT Point-to-Point Connections	Displays all of the defined MPT point-to-point connections in your network.
MPT Point-Point Actual Path	Displays the actual path for a selected MPT point-to-point connection.
Switch Name	There are two switch name fields. Each field displays the name of each of the switch endpoints of a selected MPT point-to-point connection.
Switch ID	There are two switch ID fields. Each field displays a unique identifier that is assigned to the two switch endpoints of a selected MPT point-to-point connection. (This value is the last two bytes of the switch IP address.)
Oper Info	Displays *Up* or *Down* to indicate the current operational status of a selected MPT point-to-point connection.
Hop Count	Displays the number of hops used in the path for a selected MPT point-to-point connection.
Using Defined Path	Displays one of the following values: *Yes* – Indicates that the point-to-point connection uses a user-defined path. *No* – Indicates that the path uses the point-to-point connection that was automatically defined by the Virtual Network Navigator (VNN).
Fail Reason	Displays the word *None* if no failure exists or displays the reason in the event of a failure. These fail reasons are reported by the switch to the NMS. Table 12-2 on page 12-9 displays a list of possible fail reasons.
Failed Node	Displays one of the following values: *No Failed Node* – Indicates that no failure exists. *A Node ID value* – Indicates a failure in the displayed node ID.
Failed Port	Displays one of the following values: *No failed port* – Indicates that no failure exists. *Logical Port Interface Number* – Indicates the number that identifies a logical port interface (that MPT is using to access the switch) in the event of a failure.

2. Choose Add. The Select MPT Point-to-Point Connection Endpoints dialog box appears (Figure 12-5).

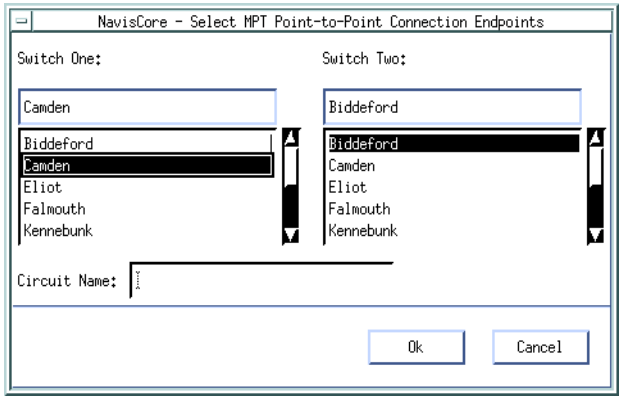

Figure 12-5. Select MPT Point-to-Point Connection Endpoints Dialog Box

3. Use the up and down arrows to select the two switches as the endpoints for this MPT Point-to-Point connection.

   A switch does not appear in the selection list if:

   - The MPT administrative value for the switch is set to Disable. See page 12-1 for more information about checking the MPT administrative value.
   - There are no IP interfaces configured for the switch.

4. Specify the Circuit Name for this connection.

5. Choose OK. The Show All MPT Point-to-Point Connections dialog box reappears and the Point-to-Point connection is included in the list of connections. The Virtual Network Navigator (VNN) automatically defines the best path for an MPT point-to-point connection. However, you can use the Define Path option to create a user-defined path. See the following section, "Defining an MPT Point-to-Point Connection Path" for details.

Defining an MPT Point-to-Point Connection Path

VNN automatically uses the best path for an MPT point-to-point connection when you add the point-to-point connection. However, you can use the Define Path option to create a user-defined path for an existing MPT point-to-point connection.

To define the path for an MPT Point-to-Point Connection:

1. Select the connection from the Set All MPT Point-to-Point Connections dialog box (Figure 12-4 on page 12-8).
2. Choose Define Path. The Set MPT Point-to-Point Define Path dialog box appears.

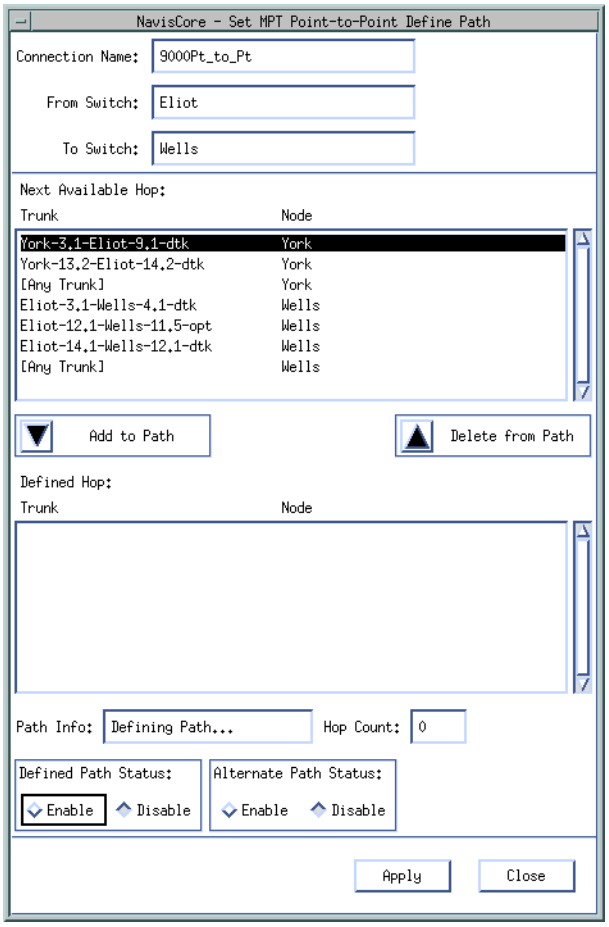

Figure 12-6. Set MPT Point-to-Point Define Path Dialog Box

3. Use the Add to Path arrow to add a hop from the Next Available Hop list to the Defined Hop list. Use the Delete From Path arrow to delete a hop from the Defined Hop list.
4. Specify the necessary field information for the Defined and Alternate Path status as shown in Table 12-3.

Table 12-3. Set MPT Point-to-Point Defined Path Field Descriptions

Field	Action/Description
Connection Name	Displays a unique name that identifies the MPT Point-to-Point connection.
From Switch	Displays a name that identifies the first switch endpoint of the MPT point-to-point connection.
To Switch	Displays a name that identifies the second switch endpoint of the point-to-point connection.
Next Available Hop	Lists the trunks that are available for use in defining the path for this point-to-point connection.
Defined Hop	Lists the trunks that you are using for this user-defined point-to-point connection.
Path Info	Displays the status of the user-defined path.
Hop Count	Displays the number of hops used in the user-defined path for this MPT point-to-point connection.
Defined Path Status	Select Disable to indicate that the path is administratively Down. Select Enable to indicate that the path is administratively Up.
Alternate Path Status	Select Disable to indicate that the alternate path is administratively Down. The alternate path is the path that VNN automatically defined before you created that user-defined path. If you select Disable and the user-defined path for the MPT point-to-point connection fails, VNN will not use the alternate path. Select Enable to indicate that the alternate path is administratively Up. If you select Enable, VNN will use the alternate path if the user-defined path fails.

Displaying the Operational Status

To display the operational status for an MPT Point-to-Point connection path:

1. Select the connection from the Set All MPT Point-to-Point Connections dialog box (Figure 12-7).
2. Choose Oper Info. The system then polls the network and updates the information in the right-hand portion of the dialog box if necessary.

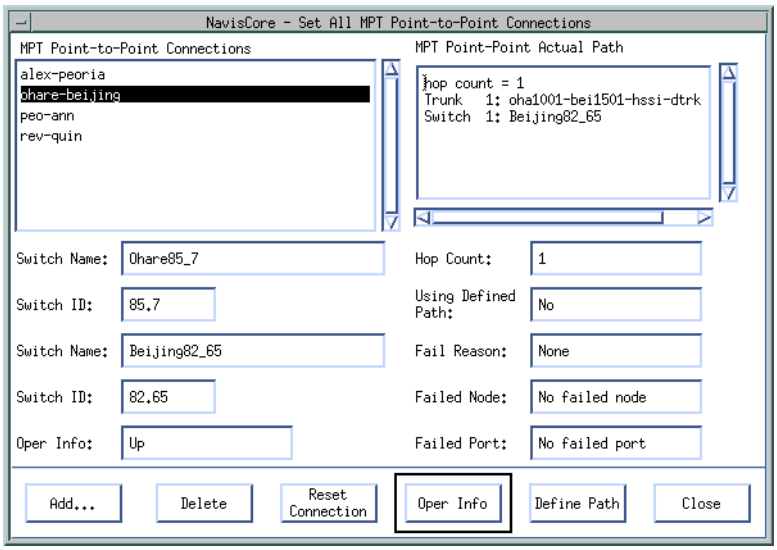

Figure 12-7. Displaying the Operational Status

See Table 12-4 for a description of each of the possible failure reasons.

Table 12-4. MPT Point-to-Point Connection Failure Reasons

Field	Action/Description
None	No failure condition exists.
Unknown	The error condition that the switch reported does not match any of the error conditions in this table.
Tpcalling	VNN detected an MPT port calling error.
Vcalling	VNN detected a circuit endpoint calling error.
Tpdead	A dead MPT port exists.
Route Lookup	A route lookup failure exists.
Confirm Timeout	The confirm timer has expired.
Path Clear	An OSPF Path clear condition exists.
Trunk Down	A trunk down condition exists.
Dead	Hello packets are no longer being received.
Grooming	A better path exists in the network.
Path Registering	Indicates a failure to register a path with OSPF.
Impure path	A shorter mixed (cell/frame) path exists.
Receiver died	The Reassembly Virtual Circuit (RVC) has died.
Invalid Trunk	An invalid trunk was found in the path.
VPC Allocate	VNN could not allocate a VPC.
VCC Allocate	VNN could not allocate a VCC.
ADD RVC	VNN could not add a reassembly virtual circuit (RVC).
NO RIDS	No more RIDs are available at this node.
FE Calling	VNN encountered an MPT Forwarding Engine (FE) calling error.
FE Dead	VNN was notified of a dead MPT Forwarding Engine (FE).
FE Active	VNN was notified of an active MPT Forwarding Engine (FE).
Invalid FE ID	An invalid Forwarding Engine (FE) ID exists at the leaf.

A

PRAM Upload

This appendix describes the uses for the Upload PRAM feature and supported IP objects. For complete details about Upload PRAM, see the *NavisCore Getting Started Guide*.

Using the Upload PRAM Command

Occasionally the switch configuration file for a specific I/O module and the configuration stored in the NMS database do not match. This situation can occur when you upgrade your switch software, use a network management product to manage the switch, or use the MIB to change a switch configuration.

If you remove an I/O Module from one switch and install this module in a second switch, you get a PRAM conflict. This happens because the module contains an unknown configuration. Do not use PRAM upload to clear this condition. Instead, use the Erase PRAM function to clear PRAM on this module; then reconfigure the module. Refer to the *NavisCore Getting Started Guide* for more information.

To resolve PRAM conflicts, use the Upload PRAM function to view the switch configuration file stored in PRAM. This enables you to compare the configuration file in the switch (PRAM) to the configuration file in the NMS database.

Using Upload PRAM After Configuring IP Objects

Upload PRAM currently does not support MPT Point-to-Point connections. All other IP objects are supported. See the *NavisCore Console Command Reference Guide* for more information.

Before you use the Upload PRAM function, review the following points:

- If you configure IP parameters using the console commands instead of using NavisCore, you need to upload the switch configuration to the NMS database. See the *NavisCore Getting Started Guide* for detailed instructions about how to upload a switch configuration file.
- You can use Upload PRAM to add objects from switch PRAM to the NMS database, as long as the objects being added do not conflict with existing objects in the database; for example, if the NMS database already contains a switch with a switch that you are adding, there would be a conflict.
- Due to the interdependency of objects with other objects in the database, *be careful* when you use Upload PRAM to delete objects from the database. In general, do not create a situation where there are dangling objects (i.e., an object without a parent) in the switch before applying Upload PRAM.

  For example, deleting a logical port without first deleting all associated individual addresses or address screens, creates dangling objects and causes a problem during the Upload PRAM process.

Glossary

A

ARP

See Address Resolution Protocol.

Address Resolution Protocol

IP services uses this protocol to resolve the lack of a hardware address when an IP address of a given destination is known but the destination hardware address (DLCI or VPI/VCI) is not.

AS

See Autonomous System.

Asynchronous Transfer Mode

A method used for transmitting voice, video, and data over high-speed LAN and WAN networks.

ATM

See Asynchronous Transfer Mode.

Autonomous System

A set of routers having a single routing policy running under a single technical administration.

B

Backward Explicit Congestion Notification

A bit in the Frame Relay header that indicates the frame has passed through a congested node from traffic traveling in the opposite direction.

Bc

See Committed Burst Size.

Be

See Excess Burst.

BECN

See Backward Explicit Congestion Notification.

BGP

See Border Gateway Protocol.

Border Gateway Protocol

A protocol that exchanges routing information between autonomous systems.

C

CIDR

See Classless Inter-Domain Routing**.**

CIR

See Committed Information Rate.

circuit

A communications channel or path between two devices.

Classless Inter-Domain Routing

An interim solution to the scarcity of IP addresses. CIDR allows routers to group routes together. This practice minimizes the quantity of information carried by core routers. CIDR requires that IP addresses and their subnet masks be written as 4 octets, separated by periods, followed by a forward slash (/) and a two-digit number that indicates the subnet mask.

Committed Burst Size

The maximum amount of data, in bits, that the network agrees to transfer under normal conditions, during a time interval Tc. Committed Burst Size is defined for each PVC.

Committed Information Rate

The rate at which the network agrees to transfer information under normal conditions. The rate is averaged over a minimum increment of time, Tc.

D

datagram

A message unit that contains source- and destination-address information, as well as the data itself, which is routed through a packet-switched network.

Data Link Connection Identifier

A 10-bit address that identifies PVCs.

DLCI

See Data Link Connection Identifier.

Distance Vector Multicast Routing Protocol

A multicast routing protocol that maintains its own unicast routing table by exchanging routing messages with DVMRP neighbors. In this technique a router that receives a multicast packet needs to know if other multicast routers to which it has connections need to receive the packet. DVMRP sends the packet to all attached routers and waits for a reply. Routers with no group members return a *prune* packet to allow routers to leave a multicast group.

DVMRP

See Distance Vector Multicast Routing Protocol.

E

Ethernet

A popular LAN protocol and cabling scheme with a transfer rate of 10 Mbps.

Excess Burst

The maximum allowed amount of uncommitted data (in bits) in excess of Bc that the network attempts to deliver during time interval Tc. In general, this data (Be) is delivered with a lower probability than Bc.

F

FECN

See Forward Explicit Congestion Notification bit.

File Transfer Protocol

A method of transferring information from one computer to another, either over a modem and telephone line, or over a network. FTP is a TCP/IP application utility.

Forward Explicit Congestion Notification bit

A bit in the Frame Relay header that indicates the frame has passed through a node that is experiencing congestion in the same direction in which the frame is traveling.

Frame Relay

A type of data transmission based on a packet-switching protocol, with transmission rates up to 2 Mbps. Frame Relay provides for bandwidth-on-demand.

FTP

See File Transfer Protocol.

I

ICMP

See Internet Control Message Protocol.

IFMP

See Ipsilon Flow Management Protocol.

IGMP

See Internet Group Multicast Protocol.

InARP

See Inverse Address Resolution Protocol.

Input/Output Adapter

A module that connects the various IOP and IOP Plus modules in a switch. IOA configurations vary according to the specific IOP module they support.

Input/Output Processor

A module in a switch that manages the lowest level of a node's trunk or user interfaces. An IOP performs physical data link and multiplexing operations on external trunks and user links.

Internet Control Message Protocol

An error-reporting protocol that works with IP.

Internet Group Multicast Protocol

An Internet protocol that is used to join or remove a host from a multicast group.

Internet Protocol

The TCP/IP session-layer protocol that regulates packet forwarding. See also Internet Control Message Protocol.

Internet Protocol address

A 32-bit address assigned to hosts using TCP/IP. The address is written as four octets separated with periods (dotted decimal format), which are made up of a network section, an optional subnet section, and a host section.

Inverse Address Resolution Protocol

IP services uses this protocol to resolve the lack of a hardware address when the destination hardware address (DLCI or VPI/VCI) of a given destination is known but the destination IP address is not.

Ipsilon Flow Management Protocol

(RFC 1953). A protocol that communicates route information between controllers in an IP switching environment.

IP

See Internet Protocol.

IP address

See Internet Protocol address.

K

Kbps

Kilobits per second.

M

MAC Address

A standardized data link layer address that is required for every port or device that connects to a LAN.

Management Information Base

The set of variables forming a database contained in a CMIP or SNMP-managed node on a network. Network management stations can fetch/store information from/to this database.

Mbps

Megabits per second.

MIB

See Management Information Base.

MOSPF

See Multicast Open Shortest Path First.

MPOA

See Multi Protocol Over ATM.

MPT

See Multipoint-to-Point-Tunneling.

Multicast Open Shortest Path First

A protocol that is designed for use with autonomous systems. MOSPF is an Open Shortest Path First routing protocol that determines the best path through the network to a specific multicast router.

Multipoint-to-Point-Tunneling

MPT allows multiple nodes to share the same circuit for transmission to a single destination. An MPT can be thought of as the inverse of the point-to-multipoint virtual circuit used to allow the sending of packets from one source to multiple destinations. Connections of this type are commonly used in multicast applications such as video distribution.

Multi Protocol Over ATM

A method for overlaying Layer 3 network routing protocols (like IP) over an ATM switched network environment. MPOA is an ATM Forum specification.

N

NavisCore

The UNIX-based graphical user interface used to configure and monitor an Ascend switch network.

NBMA

See Non-Broadcast Multi-Access.

NHRP

See Next Hop Resolution Protocol.

network address

A network layer address refers to a logical, rather than a physical network device; also called protocol address.

Next Hop Resolution Protocol

An address resolution protocol that provides a source station (host or router) with the NBMA address of the next hop to a destination station.

node

Any device such as a PC, terminal, workstation, etc., connected to a network and capable of communicating with other devices.

Non-Broadcast Multi-Access

A network to which multiple computers and devices are attached, but data is transmitted directly from one computer to another over a virtual circuit or across a switching fabric. ATM and Frame Relay are examples of NBMA media.

O

Open Shortest Path First

A routing protocol that takes into account network loading and bandwidth when routing information over the network. Incorporates least-cost routing, equal-cost routing, and load balancing.

OPTimum PVC trunk

A logical port configuration that optimizes interoperability in performance and throughput in networks where both ends are connected by Ascend switches.

OPTimum trunking

A software function that allows public data networks based on Frame Relay, SMDS, or ATM to be used as trunk connections between Ascend switches.

OSPF

See Open Shortest Path First.

P

packet

Any block of data sent over a network. Each packet contains sender, receiver, and error-control information in addition to the actual message; sometimes called payload or data bits.

Parameter Random Access Memory

The PRAM on a switch that contains the module's downloaded configuration file, and which is stored in battery backup.

Peak Cell Rate

In ATM transmission, the maximum transmission rate that cells are transmitted. Equivalent to Be for Frame Relay, PCR is measured in cells per second and converted internally to bits per second. PCR defines the shortest time period between two cells.

Permanent Virtual Circuit

A logical connection across a packet-switched network that is always in place and always available along a predetermined network path.

PIM

See Protocol Independent Multicast.

PNNI

See Private Network to Network Interface.

Point-to-Point Protocol

A protocol that provides router-to-router and host-to-network connections.

PPP

See PPP.

PRAM

See Parameter Random Access Memory.

Private Network to Network Interface

An ATM routing and signalling protocol jointly designed by member companies of the ATM Forum. It serves as an industry standard for dynamically routing scalable ATM switched Virtual Circuits (SVCs).

Protocol Independent Multicast

A multicast routing protocol that can provide scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol. PIM has two operational modes, one for densely distributed multicast groups and one for sparsely distributed multicast groups.

protocol

A set of rules governing communication between two entities or systems to provide interoperability between services and vendors. Protocols operate at different layers of the network, e.g., data link, network, and session.

PVC

See Permanent Virtual Circuit.

Q

QoS

See Quality of Service.

Quality of Service

A statistical report that specifies certain characteristics of network services, sessions, connections, or links. For example, a NavisCore statistics report describes the lost packets and round-trip delay measurements.

R

Request For Comment

A series of notes and documents available on-line that describe surveys, measurements, ideas, techniques, and observations, as well as proposed and accepted Internet protocol standards, such as Telnet and FTP.

RFC

See Request For Comment.

RFC1294

A specification documenting multi-protocol access over Frame Relay.

RIP

See Routing Information Protocol.

router

An intelligent LAN-connection device that routes packets to the correct LAN segment destination address(es). The extended LAN segments may or may not use the same protocols. Routers link LAN segments at the ISO/OSI network layer.

routing

The process of directing data from a source node to a destination node.

Routing Information Protocol

A routing protocol that maintains a list of accessible networks and calculates the lowest hop count from a particular location to a specific network.

routing protocol

A protocol that implements routing using a specific routing algorithm. Routing protocols include IGRP, OSPF, and RIP.

S

static route

A route or path that is manually entered into the routing table. Static routes take precedence over routes or paths specified by dynamic routing protocols.

subnet address

An extension of the Internet addressing scheme that allows a site to use a single Internet address for multiple physical networks.

subnet mask

A 32-bit address mask used in IP to specify a particular subnet. See also *address mask*.

Switched Virtual Circuit

A logical connection across a packet-switched network providing as-needed connections to any other node in the network.

T

TCP

See Transmission Control Protocol.

topology

The map or configuration design of a network. Physical topology refers to the location of hardware. Logical topology refers to the paths that messages take to get from one node to another.

traffic shaping

In Frame Relay, a set of rules that describes traffic flow. The sender has a mechanism to ensure that the transmission of its guaranteed packets behaves in a certain way. The network knows what kind of traffic to expect, and can monitor the behavior of the traffic.

Transmission Control Protocol

The Internet standard, transport-level protocol that provides the reliable, full duplex, stream service on which many application protocols depend.

trunk

The communications circuit between two switches.

V

VC

See Virtual Channel; Virtual Circuit.

VCI

See Virtual Circuit Identifier.

Virtual Channel

A connection between two communicating ATM networks.

Virtual Circuit

A logical circuit set up to ensure reliable communication between two network devices.

Virtual Circuit Identifier

A 16-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.

Virtual Network Navigator

The connection-oriented routing technology used in Ascend switches.

Virtual Path

A group of VCs carried between two points that provides a way to bundle traffic headed in the same direction.

Virtual Path Identifier

An 8-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.

Virtual Private Network

A network that provides dedicated bandwidth and guaranteed performance, reliability, and privacy.

VNN

See VP.

VP

See Virtual Path.

VPI

See Virtual Path Identifier.

VPN

See Virtual Private Network.

Index

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V